Time-Varying Expenditure Shares and Macroeconomic Dynamics*

Benjamín García[†] Mario Giarda[‡] Carlos Lizama[§] Damian Romero[¶]

This version: October, 2023

Abstract

We examine the impact of income heterogeneity on macroeconomic dynamics by analyzing households' expenditure decisions across different goods over the business cycle. Using Chilean transaction-level expenditure data, we observe income-dependent systematic variations in expenditure shares over the business cycle, suggesting a relevant role for non-homothetic preferences. We embed these preferences into a Heterogeneous Agent New Keynesian model and analyze their influence on the transmission of fiscal transfers. We find two novel channels associated with non-homotheticities: aggregate consumption sensitivity to income and insurance through expenditure switching. In a calibration for Chile, we find that non-homotheticities lead to substantial amplification of the effects of fiscal transfers of up to fifty percent.

JEL Codes: D14, D31, E21, E32, E62

Keywords: Consumption, Business Cycles, Non-homothetic Preferences, Heterogeneous Agents, Fiscal Policy

*The views expressed are those of the authors and do not necessarily represent the views of the Central Bank of Chile or its board members. This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in the economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities due to collaboration agreements signed with these institutions. The authors thank the comments of assistants at LACEA 2022, Universidad Católica de Chile, the Annual Conference of the Central Bank of Brazil 2023, Congress of the Catalan Economic Society 2023, Society for Economic Dynamics Congress 2023, and EEA-ESEM Congress 2023.

[†]Central Bank of Chile. Email: bgarcia@bcentral.cl

[‡]Central Bank of Chile. Email: mgiarda@bcentral.cl

[§]Central Bank of Chile. Email: clizama@bcentral.cl

[¶]Central Bank of Chile. Email: dromeroc@bcentral.cl

1 Introduction

The recent strand of the literature on New Keynesian models with heterogeneity, commonly known as HANK models (see Kaplan et al., 2018; Auclert, 2019), has highlighted the role of income fluctuations in shaping the business cycle. There has not been, however, significant work on the general equilibrium implications of households endogenously redistributing their spending across different goods, such as food, manufacturing, and services.

At the same time, multi-sector new Keynesian models (see Pasten et al., 2020 and Baqaee et al., 2021 among others) have been developed to study the effect of aggregate shocks on the sectoral allocation of resources through supply-side channels. While these models allow for the analysis of the propagation of shocks between different sectors, they offer limited insights into the role of consumption bundle decisions in sectoral allocation. This limitation arises in these models because they assume homothetic preferences, such as constant elasticity of substitution (CES), which do not generate fluctuations in expenditure shares beyond changes in relative prices.

In this paper, we aim to bridge these gaps by studying the role of income and household heterogeneity in sectoral reallocation over the business cycle. We start by showing empirically that expenditure shares across goods correlate with households' income level (Cravino and Levchenko, 2017; Cravino et al., 2020). Using survey data from Chile, we show that the share of consumption spending allocated to food ranges from 35% in the lowest decile of income to 13% in the highest decile. In contrast, the share devoted to services increases from 46% to 69% as we move from the first to the tenth decile of the income distribution.[1] These results only provide evidence of a static relationship between income and expenditure shares. For this reason, we provide evidence from credit and debit card transactions that allow us to extend the analysis into a dynamic dimension and at a monthly frequency. We do so by examining the cyclicality of expenditure shares and the dynamic impact that income shocks have on those shares. We find that expenditure shares vary over the business cycle in accordance with the income elasticities from the static analysis. In particular, the change in the share of food is negatively correlated with the change in output, while the correlation for the share of manufactured goods and services is positive. Additional analysis at the municipal level further confirms those results.[2] To further study the relationship between income and expenditure shares beyond that unconditional cyclicality, we analyze the effects of fiscal transfers on expenditure shares by using data on fiscal transfers to households from 2018 to 2021, which

[1]While this evidence is particular to Chile, the general patterns associated with differential income elasticities are a more generalized stylized fact as we show in Appendix A.2 for OECD countries. Chile is an interesting case for the analysis as it features similar patterns as the US but with more pronounced relationships between consumption shares and income.

[2]We show that these results also hold during the COVID pandemic.

include unprecedented fiscal help made by the Chilean government during the COVID-19 pandemic. We find that after a fiscal transfer, the share spent on services increases the most while the share on food declines. This result suggests that income elasticities affect the patterns of expenditure shares in the cross-section of households and generate dynamic effects on sectoral and aggregate consumption. Our findings emphasize the importance of treating preferences for different goods as an endogenous and time-varying phenomenon, not only as a switch in the relative position of households in the income distribution as generally implied by previous literature (Clayton et al., 2018; Cravino et al., 2020).

Our empirical evidence implies that expenditure shares are heterogeneous across the population, are time-varying, and are endogenous to the business cycle. These facts motivate us to build a HANK model, embedded with NH preferences, in which there exists heterogeneity in income and consumption due to incomplete markets, which are crucial elements to match the empirical facts. In our setup, households with different income levels display different expenditure shares, different price indices, and different average income elasticities.[3] In our framework, the household's average income elasticity is a crucial element since it determines the sensitivity of the household's consumption composition to income shocks and then the sensitivity of the economy as a whole.

This novel source of heterogeneity interacts with the traditional heterogeneity in MPCs due to the occasionally binding financial constraints that HANK models have. In particular, our model implies that poorer households feature lower average income elasticities and larger MPCs than wealthy households. At a household level, NH preferences provide insurance by allowing an optimal bundle rebalancing after an income shock, leading to lower consumption variability. At an aggregate level, however, if the economy-wide average income elasticity increases sufficiently in response to an income shock, households become, on average, more sensitive to income changes, amplifying the aggregate consumption response. These channels are, respectively, *insurance* and *average elasticities*. We formally decompose the effects on consumption between these two channels following Patterson (2019). We show that the aggregate average elasticity, if procyclical, amplifies the effects of income shocks (such as fiscal transfers) given an average MPC, while the covariance between household elasticities and their MPCs generates further amplification or dampening of shocks, depending on the sign of this covariance. Ultimately, the response to shocks depends on the relative weight of these two channels.

We show that in general equilibrium, a positive relationship between the income elasticities and output tends to appear, suggesting that whenever the elasticity responds more strongly, there is a higher amplification of fiscal transfer shocks. We provide several robustness checks by studying the role of

[3] We define the average income elasticity of a household as the mean of the income elasticities of the goods in its consumption bundle, where the income elasticity of a good represents the change in the participation of the good in the household's consumption bundle given a change in income.

monetary policy, the financing of fiscal transfers, and the persistence of shocks. We show that our results are robust to the model in place and that the contribution of income elasticities to the amplification of fiscal transfer shocks can be sizable.

This paper contributes to the literature as follows. First, we empirically show with granular data at a monthly frequency that changes in income, both aggregate and idiosyncratic, lead to adjustments in the household's relative demand for different goods. This confirms the presence of NH mechanisms that drive fluctuations in expenditure shares. Second, we show that the variability of the household's average income elasticity across the income distribution has significant effects on consumption and GDP dynamics as these elasticities can act as amplifiers (through higher sensitivity) or dampeners (through insurance by bundle switching) of the effects of shocks. We demonstrate the relevance of expenditure switching conditional on fiscal transfer shocks. There are two papers related to ours. First, Ferrante et al. (forthcoming) show that shifts in expenditure composition have relevant effects on inflation and aggregate dynamics. They estimate a DSGE model with a "taste shifter shock" that exogenously switches the relative demand for the different goods (services and goods). We extend their findings by including an endogenous bundle switcher with empirical validity that arises due to NH preferences. Second, our paper has similarities with Schaab and Tan (2023) who embed NH preferences into a HANK with input-output linkages and heterogeneity in income. The contribution of our paper with respect to theirs is that we isolate and explain the time-varying bundles' mechanism in their model while also–and more importantly–providing empirical evidence of the existence of these mechanisms.

The remainder of the paper is organized as follows. Section 2 presents the empirical evidence for expenditure heterogeneity across the income distribution and for the dynamic responses of expenditure shares in the business cycle. In section 3, we describe our HANK model with non-homothetic preferences. In section 4, we study the role of these preferences in a HANK that features a distribution of MPCs. We show the quantitative exercises in section 5. And finally, section 6 concludes.

2 Empirical Evidence

In this section, we document three empirical facts about the composition of expenditures of Chilean households and their dynamic behavior. First, we analyze data from the expenditure surveys to show that the expenditures' composition varies along the income distribution. That implies that different groups of goods have different income elasticities. Second, we use data from credit and debit card transactions to show that in the business cycle, the expenditure shares fluctuate in patterns consistent with the static elasticities presented by Fact 1. Finally, we show that even after controlling for business cycle variables, consumption shares respond to fiscal transfers according to the income elasticities we show before.

2.1 Data

Expenditures at the Household Level. First, for the analysis of consumption expenditures at the household level from a static perspective, we rely on the Consumption Expenditure Survey (*Encuesta de Presupuestos Familiares* in Spanish, or EPF hereafter). It is a cross-sectional survey that provides information about expenditures and household characteristics (demographics and income). The main goal of EPF is to serve as the base for constructing the Consumer Price Index (CPI).[4]

While this is the primary source of information for consumption expenditures in Chile, using this dataset carries two main issues. First, unlike expenditure surveys in the US, such as the Consumer Expenditure Survey (CEX), we cannot follow individuals or households over time but only construct pseudo-panels to analyze expenditure patterns. Therefore, we can not control for unobservable characteristics that might be important to understand consumption patterns. Second, the dataset is not taken at business-cycle frequency (i.e., every quarter) or regular intervals. In particular, the survey's most recent waves were taken in 1996, 2007, 2013, and 2017. Because of the reasons above, we base our analysis on the most recent wave of the data, EPF VIII for 2017, to document the heterogeneity in consumption expenditures and provide robustness of our results in the appendix for previous waves of the data.

EPF VIII provides information on different sources of income, such as labor income, rents from assets and real estate, and imputed rentals for homeowners. Following the literature (Cravino et al., 2020), our baseline classification uses total disposable income per capita at the household level, which is the sum of all income sources net of transfers. We then classify each household according to its percentile in this distribution. In Figure A.1 in the appendix, we show that the general patterns presented below are similar if we use the distribution of labor income instead of total income. Also, using a more narrow classification, such as deciles, does not change the general picture presented in this section.

Expenditure in EPF VIII is structured into five different levels of aggregation. From more to less disaggregated, these groups are defined by 1,186 products, 285 subclasses, 126 classes, 59 groups, and 12 divisions. We aggregate the 12 divisions to generate three categories: food and beverages, manufactured goods, and services.[5] Table A.1 gives more details about their components and the expenditure shares across selected percentiles of the income distribution.

Expenditures at the Municipality Level. We use monthly data on credit and debit card transactions as a proxy of expenditures from Transbank, a private firm that processes most of Chile's credit and debit transactions. This data allows us to capture the expenditure behavior of different units of individuals

[4]For this analysis, we follow Romero (2022) closely.

[5]We prefer to work with this more coarse classification rather than a granular one in order to (i) use data on prices to estimate key parameters of the model and (ii) be able to compare expenditure patterns over time. See details below.

over time to study their cyclical behavior.

However, using this data as a proxy of aggregate expenditures has three shortcomings. First, this data distinguishes between in-person and online purchases. We use the former, as online purchases are harder to associate with the buyer's residence and limit us to the panel data analysis below. Also, although we can observe the firm and place of the transactions, we do not know the individual who made the purchase. Due to these restrictions, we restrict ourselves to conducting our analysis at the municipality level.[6]

Second, one might be concerned about how well a municipality represents the income distribution of individuals. Figure A.3 of Appendix A.3 shows that municipalities do it well in representing income inequality in Chile. Represented by the mode, there is almost a one-to-one relationship between the quintile of the individuals in each municipality and the quintile the municipality belongs.

Third, this data only considers card transactions and hence represents only a fraction of the aggregate consumption in the economy (for example, by excluding cash purchases). Therefore, this data might represent financial development rather than expenditures at the aggregate level. To alleviate this concern, Figure A.4a shows that the Transbank data captures fluctuations in aggregate consumption quite well. The correlation of the growth of these series with consumption of National Accounts is 0.86 for the period before the pandemic and about 0.93, including the pandemic. The series by services and goods (manufactures plus services) also matches National Accounts well, as we show in Figures A.4b and A.4c in Appendix A.3.

Fiscal Support. From the Ministry of Social Security and the Chilean Pensions Supervisor, we obtain the total amount of different fiscal transfer programs, allowing us to control for conditional correlation between income and expenditure shares. We have information about direct fiscal transfers from 2018 to 2022 in Chile. The measures considered include the policies undertaken during the COVID-19 pandemic and several other programs starting in 2018.[7]

While this data is available at the individual level, we aggregate it at a municipality level for our empirical analysis as it allows us to draw a direct comparison with our measures for time-varying expenditures from card transactions, only available up to that level of aggregation. Figure 1 shows the programs' size. As the figure shows, there was sizeable fiscal support during COVID-19, especially in

[6]The geographical approach is used, for instance, by Mian and Sufi (2009) and Mian et al. (2013) to study the effects of wealth on consumption. This approach is also extensively discussed by Guren et al. (2020) to disentangle general equilibrium from the partial equilibrium effects of these estimates.

[7]We have eleven programs featuring different sizes, timings, cyclicality, and progressivity. Those programs are (i) Family help check; (ii) Family base check; (iii) Christmas COVID check; (iv) School homework check; (v) Child homework check; (vi) COVID emergency check; (vii) Protection check; (viii) Emergency Income COVID; (ix) Emergency COVID 2020; (x) Guaranteed Minimum Income; and (xi) Universal COVID check.

2021. For completeness, we compare those measures with the additional liquidity that households got through three projects, allowing them to withdraw at most ten percent of individual pension funds three times (P.F. Withdrawals in the plot). Later on, in our empirical analysis, we focus on the impact of Fiscal Transfers; however, we control for the effect of those withdrawals to properly account for the impact of fiscal support over expenditure shares. In particular, we study the effect of the policies *per unit* of additional liquidity provided to households, conditional on pension fund withdrawals and fluctuations at a country level.

FIGURE 1: Fiscal Support in Chile

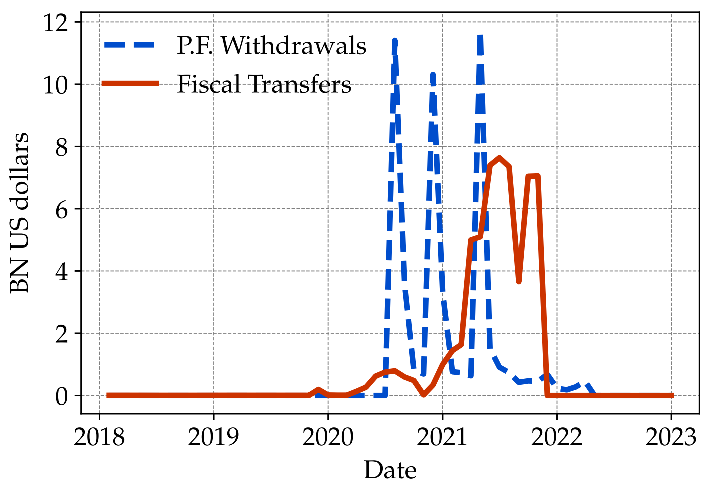

NOTES: This figure compares the evolution of total fiscal support to households over time with information over pension fund withdrawals in billions of dollars.

Aggregate Data. Finally, we use monthly data on industrial production (IMACEC) and sectoral price indexes from the Central Bank of Chile to control for business cycle covariates, as well as relative prices that might affect expenditure patterns.

2.2 Fact 1: Expenditure Shares Vary Along the Income Distribution

Figure 2 presents our first observation that motivates the theoretical and empirical analysis in the following sections: expenditure shares vary along the income distribution. Each panel presents the expenditure share for the three groups of goods (food and beverages, manufactured goods and services) relative to the total expenditures of each percentile. The figure reveals clear patterns of consumption across the income distribution. First, low-income households spent relatively more (35 percent) than wealthy households (13 percent) on food and beverages. In between these two points, there is a monotone decreasing pattern. Second, even though manufactured goods seem to have a U-shaped form, the pattern of expenditures is

FIGURE 2: Aggregate consumption patterns across income distribution

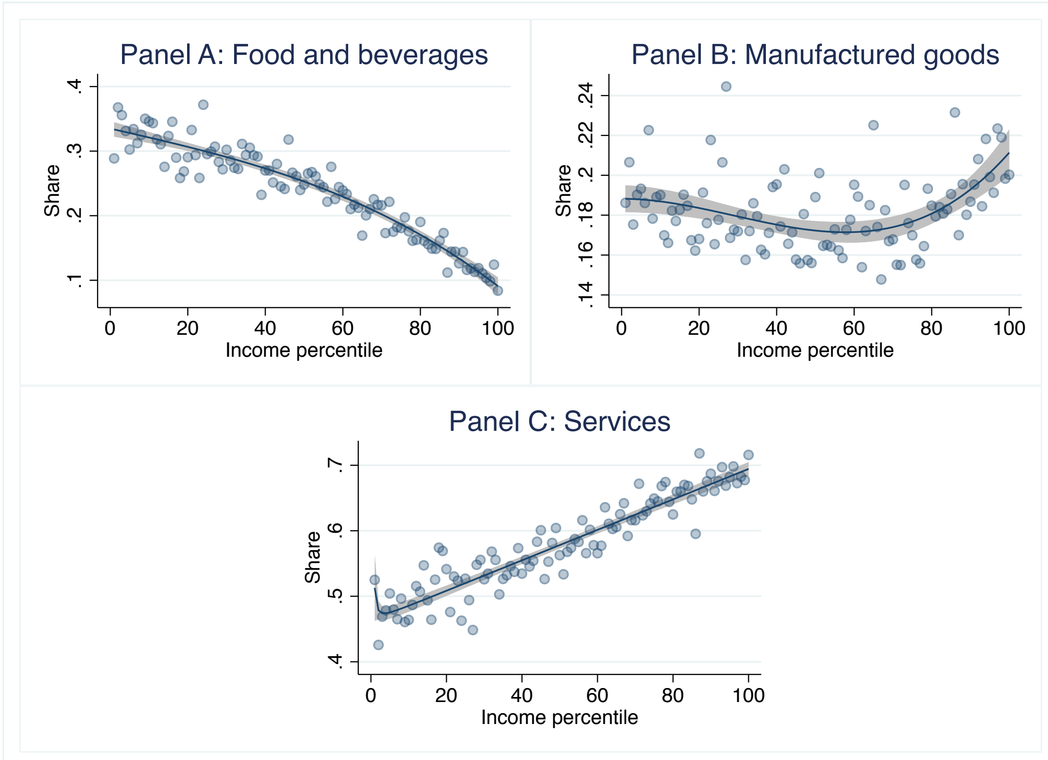

NOTES: This figure presents the share in consumption expenditures for households in each percentile of the income distribution, considering three aggregate expenditure groups. On every panel, each dot corresponds to a percentile. The solid line corresponds to the local polynomial fit. The grey area denotes 95% confidence interval.

relatively constant across the distribution. For example, the lowest decile spent 19 percent on these goods, the same as the top decile and the median household.

Interestingly, note that the expenditure pattern on these goods is more volatile than in the case of food and services. In addition, note that the classification of these goods also considers "industry goods" such as utilities (water, electricity, gas, and other fuels), which tend to have more volatile prices, which can be behind these patterns. Finally, even though services are an essential category for all households in the distribution, there is a clear increasing pattern in expenditures. While the lowest decile spent around 46 percent on these categories, the top decile spent 69 percent. All these patterns are robust to consider a more disaggregated level of the data, as well as other points in the income distribution (see Table A.1).

As mentioned at the beginning of this section, we cannot follow the same household to verify how stable these patterns are over time or across the cycle. To partially address this concern, we compare expenditures in these three broad groups of goods for different waves of the EPF survey. We present this evidence in Figure A.2, which compares the baseline distribution of expenditures in EPF VIII with those in waves V (1996), VI (2013), and VII (2013). The figure reveals that the decreasing (increasing) expenditure pattern for food (services) across the income distribution is stable over time, with a high correlation for both measures (above 80 percent). Note that the level of these expenditure shares is also quite similar. Most of the differences, however, come from the comparison between EPF V of 1996 and the baseline wave, EPF VIII of 2017, which can be due to the time difference between these two surveys and the significant changes in technological developments, as well as the changes in income over these 20 years. What is certainly different is the expenditure share in manufactured goods, which is positively but weakly correlated across surveys. This category most likely has changed its components over time, and the prices of the individual goods in each sub-category.

Income-elasticity Interpretation. The evidence presented in this section suggests that expenditure shares are not constant along the income distribution. In particular, while a negative income elasticity drives expenditures in food, expenditures in manufactures are slightly positive income-elastic, while services show a strictly positive income elasticity.

2.3 Fact 2: Expenditure Shares' Correlate with Income Elasticity over the Business Cycle

Fact 1 shows there is a static correlation between expenditure shares and income. Up to this point, this might represent an economy in which there are time-invariant differences in preferences across the income distribution (Clayton et al., 2018; Cravino et al., 2020). In this section, we show that expenditure shares fluctuate over the business cycle, both at an aggregate and a disaggregated level, consistent with the income elasticities presented in Fact 1.

The Cyclicality of Expenditure Shares at the Aggregate Level. Defining s_{jt} as the expenditure share on category j (food, manufactures, and services) in period t at the aggregate level, Figure 3 presents the relationship of the different expenditure shares (in the vertical axis) and the log-change of industrial production (in the horizontal axis) for the period January 2015 to December 2019.[8] The key observation from Figure 3 is that the composition of expenditures is time-varying. Therefore, in the cycle, consumers switch their expenditure from some goods to others. In particular, we find that during booms in the Chilean economy, consumers substitute food for manufactures and, more strongly, for services. Hence, there is a significant reallocation of spending between types of goods. The previous finding is consistent with Fact 1 while dynamically. In the business cycle, when aggregate income increases, the economy switches from low income-elastic goods to high income-elastic goods.

[8] Expenditure shares and industrial production are seasonally adjusted. We exclude the periods after COVID for expositional purposes. Figure A.5 in the Appendix uses all the available data and presents similar patterns.

FIGURE 3: Ciclicality of Aggregate Expenditure Shares

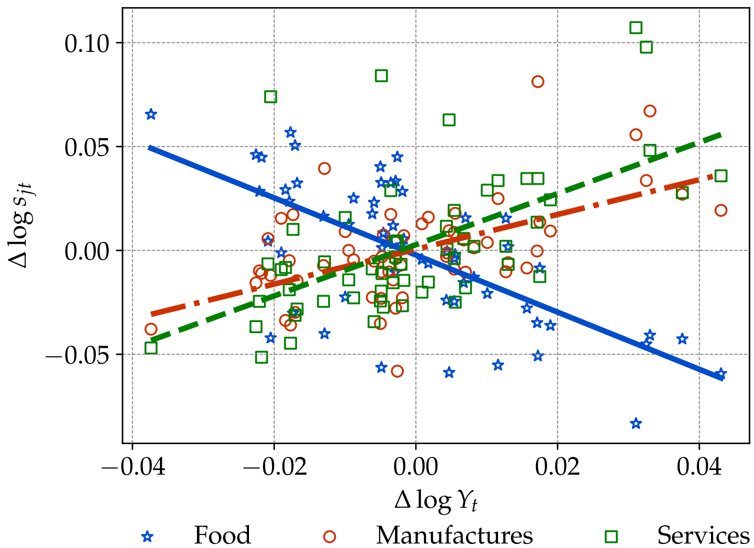

NOTES: This figure presents the unconditional correlation between the business cycle of sectoral expenditure shares and aggregate production index.

Note, however, that all these results might be driven by variations in relative prices affecting the expenditure patterns over the cycle. To investigate this possibility, we run the following regression

$$\Delta \log s_{jt} = \alpha_j + \beta_j \Delta \log Y_t + \delta_j(\pi_{jt} - \pi_t) + \varepsilon_{jt}, \tag{1}$$

where $\pi_{jt} - \pi_t$ denotes the inflation rate of category j relative to the aggregate inflation rate.

As shown in Table 1, the cyclicality of expenditure shares by category is consistent with their income elasticities. This finding is robust to controlling for relative inflation (to control for relative prices), and it is present both in the pre-and post-COVID periods.

TABLE 1: Cyclicality of Aggregate Consumption Shares

	Panel A: 2015-2022			Panel B: 2015-2019		
	Food	Manufactures	Services	Food	Manufactures	Services
$\Delta\log(Y_t)$	-1.754***	0.868***	2.311***	-1.499***	0.840***	1.157***
	(0.209)	(0.182)	(0.259)	(0.234)	(0.215)	(0.248)
$\pi_{jt} - \pi_t$	-0.416	0.747	3.532	-1.516	0.433	0.932
	(1.000)	(1.283)	(2.563)	(1.048)	(1.481)	(2.530)
N	86	86	86	50	50	50
Adj. R^2	0.457	0.405	0.503	0.463	0.166	0.340

NOTES: This table presents the aggregate relationship between expenditure shares and income. $\Delta\log(Y_t)$ denotes the monthly change in the industrial production index. π_{jt} is the yearly inflation rate of category j, while π_t is the yearly aggregate inflation rate. Each regression controls for up to nine lags of inflation differentials (not reported). Standard errors reported in parenthesis. *, ** and *** denote statistical significance at the 1, 5 and 10% levels, respectively.

The Cyclicality of Expenditure Shares at the Municipality Level. We can further exploit our data in a more granular setting. Specifically, we investigate the cyclicality of expenditure shares at the municipal level, aiming to account for potential heterogeneities that could impact the results. Defining s_{ijt} as the expenditure share in category j in municipality i in period t, we run the following regression

$$\Delta\log s_{ijt} = \alpha_j + \beta_j \Delta\log Y_t + \delta_j(\pi_{jt} - \pi_t) + \gamma\Delta\text{Mob}_{it} + \lambda_i + \varepsilon_{jt}, . \tag{2}$$

Note that Eq. (2) is a panel regression analogous to Eq. (1), which nows considers municipality fixed effects (λ_i). To avoid that all our results are driven by lockdowns during the COVID period, we also include a mobility index (Mob_{it}) capturing the time-varying heterogeneity in lockdowns across municipalities (positive changes denote less mobility restrictions).

Table 2 confirms our previous findings, showing a negative relationship between the expenditure share in food and a positive relationship with manufactures and services. Moreover, these relationships remain robust even after controlling for relative inflation and the mobility index. At this level of analysis, we find a statistically significant negative relationship between relative prices and food shares.

In contrast, the relationship between prices and the shares in manufacturing and services is not statistically different from zero. That evidence indicates the presence of non-homothetic effects, with income represented by industrial production, playing a more significant role in driving fluctuations in expenditure shares than prices. This finding aligns with the patterns observed in the aggregate data and supports the relationship between the cyclicality of expenditure shares and income elasticities.

However, including COVID-era data distorts the previous results, primarily due to the impact of lockdown measures that heavily affected the services sector. Incorporating the period after January 2020 leads to a 28 percent decrease in the cyclicality of the share of services and a 44 percent decrease in the cyclicality of food shares while maintaining the expected signs and order. These findings highlight the substantial influence of lockdowns on consumption composition, as evident from the effects of the mobility index (Mob_{it}) on the shares, which exhibit similar patterns to $\Delta \log(Y_t)$ as expected, while nonsignificant.

TABLE 2: Cyclicality of Consumption Shares at a Municipality Level

	Panel A: 2015-2022			Panel B: 2015-2019		
	Food	Manufactures	Services	Food	Manufactures	Services
$\Delta \log(Y_t)$	-0.752***	0.662**	0.937**	-1.338***	1.022**	1.300**
	(0.217)	(0.236)	(0.343)	(0.221)	(0.307)	(0.415)
$\pi_{jt} - \pi_t$	-0.897	0.833	4.588	-2.710**	0.791	1.968
	(1.008)	(1.299)	(2.558)	(1.010)	(2.107)	(3.371)
ΔMob_{it}	-0.0166	0.0145	0.0447			
	(0.0149)	(0.00826)	(0.0257)			
N	28397	27834	28478	16848	16513	16870
Adj. R^2	-0.003	-0.001	-0.001	-0.001	-0.006	-0.010

NOTES: This table presents the disaggregate relationship (at the municipality level) between expenditure shares and income. $\Delta \log(Y_t)$ denotes the monthly change in the industrial production index. π_{jt} is the yearly inflation rate of category j, while π_t is the yearly aggregate inflation rate. Mob_{it} is a mobility index capturing the time-varying heterogeneity in lockdowns across municipalities (higher values denote less mobility restrictions). Standard errors reported in parenthesis. *, ** and *** denote statistical significance at the 1, 5 and 10% levels, respectively.

2.4 Fact 3: Expenditure Shares Respond to Income Shocks in Accordance with their Income Elasticities

Facts 1 and 2 present evidence regarding the importance of income in determining expenditure shares statically and dynamically. Up to this point, it might be the case that preferences are driven by exogenous, heterogeneous, and time-varying taste shifters, as assumed by Ferrante et al. (forthcoming). This section shows that this is not the case and that expenditure shares respond to income shocks. Moreover, we show that those responses are consistent with the income elasticities presented in previous sections.

For this purpose, we study how fiscal transfers, as a proxy for income shocks, affect expenditure shares in 2018-2022. Our empirical approach follows Misra and Surico (2014), who estimate the effects of 2001 and 2008's rebates in the United States using the Consumer Expenditure Survey.[9] To be able to analyze not only the contemporaneous response of consumption to fiscal transfers shocks but also their dynamic effects, we estimate the following Local Projection like regression

$$\log s_{ijt+h} - \log s_{ijt-1} = \alpha_{jh} + \beta_{jh} \log T_{it} + \Gamma'_{jh}\mathbf{X}_{it} + \lambda_i + \lambda_t + \varepsilon_{jit+h}, \tag{3}$$

for $h = 0, \ldots, H$ and $j \in \{f, m, s\}$. Here $\log s_{ijt+h} - \log s_{ijt-1}$ denotes the log-change in expenditure shares of good j for municipality i between periods $t-1$ and $t+h$, T_{it} is the total amount of transfers received by municipality i in period t and $\mathbf{X}_{it}$ includes twelve lags of fiscal transfers and pension funds withdrawals.[10] Finally, λ_i and λ_t denote municipality and date fixed effects. The coefficient of interest is β_{jh}, which captures the cumulative response in expenditure shares up $t+h$ periods after fiscal support in period t. Importantly, we expect λ_t to absorb any variation at a country level, so the sequence of β_{jh}'s capture variation in expenditure shares beyond the ones generated by output or other macroeconomic variables.

A few remarks are in order. First, these responses do not correspond to consumption responses to exogenous fiscal transfers, even after controlling for other variables, since it may be that consumption decisions anticipate the rise in transfers. We do not claim causal identification between transfers and consumption. Second, the responses of *total* expenditures to these transfers are positive and persistent. García et al. (2023) studies this question, and we refer the reader to that work.

We show in Figure 4 the results of estimating Eq. (3), where we plot the evolution of expenditure shares after fiscal transfers. Several results are worth commenting on. First, transfers have unequal effects on expenditure shares. A rise in transfers at a municipal level raises expenditures in services disproportionately more than manufactures and food in the first six months after the shock. Then, the effect remains positive but similar to manufacturing, which becomes positive and significantly different

[9] Misra and Surico (2014) further studies the heterogeneous effects of those rebates following Johnson et al. (2006) and Parker et al. (2013). Fuster et al. (2020) also uses a similar approach, taking surveys from experiments to study the effects on consumption of raising households' income.

[10] In appendix A.6 we show robustness for this result considering different lags and controlling for the mobility index that varies along municipalities and over time. We find no significant differences to what we find in the baseline specification. We also include pension funds withdrawals in the regressors because they significantly transferred resources from illiquid accounts to households. García et al. (2023) study the differential effects of fiscal transfers and withdrawals in detail. We abstract from the latter because they coincided with strict lockdowns, mainly affecting services.

FIGURE 4: Response of Consumption Shares to Fiscal Transfers

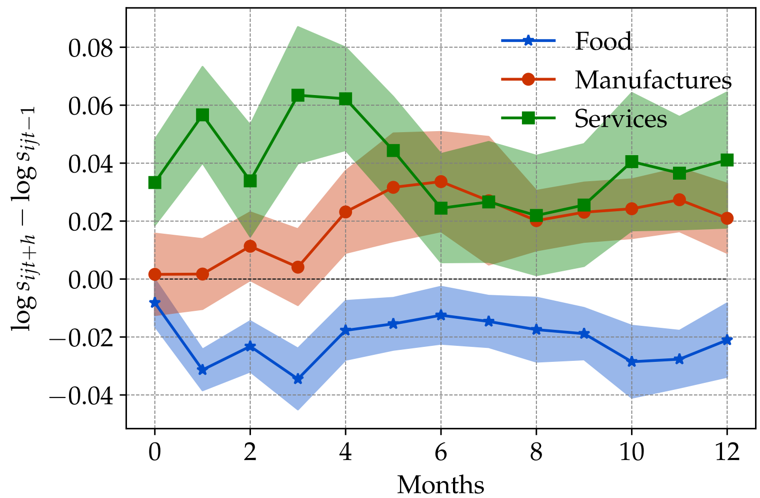

NOTES: This figure presents the dynamic response of expenditure shares to fiscal transfers (β_{jh} coefficients in. Eq. 3).

from zero four months after the shock. This effect is persistent and lasts for at least twelve periods. As expected, to the rise in the shares in manufacturing and services, a fall in foods expenditure share followed, reaching its lower four months after the shock. Thus, we find that the response to income shocks in the form of fiscal transfers is also consistent with the income elasticities found above: the share in services and manufacturing must increase while the food share must go down in response to positive income shocks. Importantly, all these results are not driven by lockdowns associated with the COVID-19 period: as we show in Figure A.6 of Appendix A.4, all our results are robust to controlling for mobility at the municipality level.

In conclusion, we observed heterogeneity in expenditure shares of different goods relative to income. Services show an increasing relationship with income, manufacturing displays a flat relationship, while food exhibits a decreasing relationship. These findings suggest that food has a negative income elasticity, manufacturing has a non-negative elasticity, and services have a positive income elasticity. Additionally, analyzing monthly credit and debit card transactions provides evidence that these patterns also hold in the evolution of consumption shares over time. Specifically, as the economy experiences growth, expenditure patterns shift from food to manufacturing and services. Furthermore, we observe these shifts in expenditure patterns also in response to income shocks in the form of fiscal transfers. What is crucial for us is that these changes in expenditures are heterogeneous, time-varying, and endogenous to the business cycle.

Motivated by this evidence, in the next section, we propose a Heterogeneous-Agent New Keynesian model with three sectors and non-homothetic preferences (NH) to provide a theoretical explanation and study the consequences of these findings.

3 A Model of Consumption with Heterogeneous and Time-Varying Expenditure Shares

This section presents a Heterogeneous-Agent New Keynesian Model (HANK) with non-homothetic (NH) preferences. To keep the exposition as simple as possible, we describe in detail the problem of households, which is the novel contribution of this paper, and briefly summarize the rest of the economy. Appendix B presents the rest of the model in more detail.

3.1 Households

We assume there is a measure of one of households that differ in wealth b and productivity z. The b-dimension is the endogenous state, and the z-dimension is the exogenous time-varying idiosyncratic shock. We identify the type of households by i, which is associated with the tuple (b, z).

Households derive utility from consumption and disutility from labor. We assume there is a continuum of $g \in (0, 1)$ labor tasks that each household (b, z) can execute. Hence, household i maximizes its lifetime utility, time-discounted at a factor $0 < \beta < 1$, given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(c_{it}, \{n_{it}^g\}_{g=0}^1\right). \tag{4}$$

Following Galí (2011), we assume a separable utility function of the form:

$$u(c_{it}, \{n_{it}^g\}_{j=0}^1) = \frac{(c_{it})^{1-\gamma}}{1-\gamma} - \chi \frac{\int_0^1 (n_{it}^g)^{1+\varphi} dg}{1+\varphi},$$

where γ is the inverse of the intertemporal elasticity of substitution, χ is the parameter of the disutility of labor, and φ is the inverse of the Frisch elasticity of the labor supply. c_{it} is total consumption and n_{it}^g is hours supplied by workers from household i to the task g. Workers do not choose their labor supply directly due to labor market frictions. The labor supply is determined by a union that represents labor supplied by households to task g, to maximize the average welfare of all households providing those services. We first show the individual problem and in the next subsection, the wage-setting process faced by the union. The implication of that assumption is that $n_{it}^g = n_t, \ \forall \ i$. Consequently, nominal wages are common across households and denoted by W_t.

Households maximize (4) subject to the following budget constraint

$$E_{it} + B_{it+1} = (1 + i_t)B_{it} + (1 - \tau_t^w)W_t n_t z_t + T_{it} + D_{it}, \tag{5}$$

where $E_{it} \equiv P_{it}c_{it}$ denotes total expenditures that are also given by

$$E_{it} = \sum_{j=1}^{J} P_{jt}c_{ijt}. \tag{6}$$

In our setting, households consume a bundle of J goods facing a (common) price P_{jt} and consuming a quantity c_{ijt}. We denote by P_{it} the price index at the household level, while c_{it} is the total household consumption i. As we explain below, our non-homothetic preferences generate household-level price indices, which imply that each household faces a different cost of living, despite observing the same sectoral prices.

Households save in a liquid and risk-free asset B_{it} that returns a nominal interest rate i_t, receive post-tax income from labor $(1-\tau_t^w)W_t n_t z_t$, where $W_t n_t$ is a "common" labor income and z_t is an individual idiosyncratic shock which generates a non-degenerate distribution of income. Households receive a fiscal transfer T_{it} and pay proportional labor income taxes where τ_{it}^w is the tax rate common to all households. Finally, total firms' dividends can be distributed unevenly among the different households, denoted by D_{it}.

The households' optimization problem can be split into an intra-temporal problem, which describes how households distribute consumption among the different goods (which in our setup are nontrivial), and the inter-temporal problem, which is also affected by allocating resources to the different goods. We study both in turn.

Intratemporal Problem–The Role of Non-homothetic Preferences. Households derive utility from the consumption of the J different goods in the economy (for our quantitative exercises, we can think about the consumption basket as composed of food, manufactures, and services). The consumption aggregator of every household is denoted by c_{it} and takes the form of an implicitly additive non-homothetic CES function. It defines total consumption in period t by[11]

$$1 = \sum_{j=1}^{J} \left(\omega_j \left(c_{it}\right)^{\epsilon_j}\right)^{\frac{1}{\sigma}} \left(c_{ijt}\right)^{\frac{\sigma-1}{\sigma}}, \tag{7}$$

where c_{ijt} denotes the consumption of good j by household i in period t, ω_j is a taste parameter for good j (common across households), σ is the constant elasticity of substitution between sectoral goods, and ϵ_j is the constant elasticity of consumption of sectoral good j with respect to the consumption index c_{it} that allows preferences to be non-homothetic. This latter parameter is also understood as the income elasticity with respect to individual consumption goods.

[11] Comin et al. (2021) shows that with these preferences, the intertemporal and intertemporal allocation problems can be separated, as in the case of standard CES utility function.

These preferences were introduced by Hanoch (1975) and recently used in the macroeconomics literature by Comin et al. (2021), Cravino and Sotelo (2019), Matsuyama (2019) and Redding and Weinstein (2019), among others, mostly to analyze trade patterns and structural change. Note that we recover the standard homothetic CES specification in the particular case of $\epsilon_j = 1 - \sigma$ for every j.

Given a level of total expenditures $E_{it} = P_{it}c_{it}$, the intratemporal cost minimization optimization problem derives the following conditions

$$c_{ijt} = \omega_j \left(\frac{P_{jt}}{P_{it}}\right)^{-\sigma} (c_{it})^{\epsilon_j+\sigma} \tag{8}$$

$$s_{ijt} \equiv \frac{P_{jt}c_{ijt}}{E_{it}} = \omega_j \left(\frac{P_{jt}}{P_{it}}\right)^{1-\sigma} (c_{it})^{\epsilon_j-(1-\sigma)}, \tag{9}$$

where s_{ijt} is the expenditure share of household i in good j. The household-specific CPI corresponds to the price index that equalizes $P_{it}c_{it} = \sum_{j=1}^{J} P_{jt}c_{ijt}$ and is defined by

$$P_{it} = \left[\sum_{j=1}^{J} (\omega_j P_{jt}^{1-\sigma})^{\vartheta_j} (s_{ijt}E_{it}^{1-\sigma})^{1-\vartheta_j}\right]^{\frac{1}{1-\sigma}}, \tag{10}$$

with $\vartheta_j \equiv (1-\sigma)/\epsilon_j$. From these expressions is clear that both the expenditure share and the CPI of each household (characterized by different levels of income and access to financial markets) depend on the level of consumption/expenditures in every period. Note again that in the case of homothetic preferences ($\epsilon = 1 - \sigma$), the expenditure shares do not depend on the level of consumption, and the CPI is common across households because it only depends on observed prices and not on the level of consumption itself (i.e., $\vartheta_j = 1$ holds). At the same time, with non-homothetic preferences, the demand for each good nonlinearly depends on total consumption through the good-specific income elasticity, ϵ_j. For future reference, denote $\bar{\epsilon}_{it} \equiv \sum_{j=1}^{J} s_{ijt}\epsilon_j$ as the average (expenditure-weighted) income elasticity. As we will see next, this object plays a crucial role in determining *intertemporal* consumption.

A visual inspection of the individual CPI in Eq. (10) suggests that the CPI is an increasing function of total consumption c_{it}. This implies that total expenditures increase non-linearly with respect to total consumption. This reflects that households switch consumption from goods with low-income elasticity to goods with high-income elasticity when they can spend one extra dollar. This is one of the crucial distinctions between a model with homothetic preferences and one with non-homothetic preferences.

Intertemporal Problem. The problem of the household can be written in the following recursive formulation, where we drop the time index and define next-period variables with a prime (recalling $i = (b, z)$, $i' = (b', z')$)

$$V(b,z) = \max_{b',c} u(c_t(b,z), n_t) + \beta \mathbb{E}_t V(b', z'), \tag{11}$$

subject to

$$\begin{aligned} p(b,z)c(b,z) + b'(b,z) &= (1+r)b + wNz + T(b,z) - \tau(b,z) + D(b,z) \\ b'(b,z) &\geq 0 \\ 1 &= \sum_{j=1}^{J} \left(\omega_j (c(b,z)^{\epsilon_j}\right)^{\frac{1}{\sigma}} c_j(b,z)^{\frac{\sigma-1}{\sigma}}. \end{aligned}$$

Note that the budget constraint of this problem is now defined in terms of the numeraire of the economy, which is aggregate consumption (see details below). Therefore, all lower case nominal variables (such as prices and bonds) are expressed relative to that numeraire and $1+r = (1+i)/(1+\pi)$, where π is the aggregate rate of inflation.

The first order conditions of problem (11) are:

$$u_c(c(b,z)) = \lambda(b,z) p(b,z) \frac{\epsilon(b,z)}{1-\sigma} \tag{12}$$

$$\lambda(b,z) = \beta \mathbb{E}_t V_b(b', z') + \mu(b,z) \tag{13}$$

$$V_b(b,z) = \lambda(b,z)(1+r) \tag{14}$$

where $\mu(b,z)$ is the Lagrange multiplier of the financial constraint, $b'(b,z) \geq 0$, and $\lambda(b,z)$ is the Lagrange multiplier of the budget constraint. Also, $\frac{p(b,z)\overline{\epsilon}(b,z)}{1-\sigma} \equiv \mathcal{E}(b,z)$ is individual price times the elasticity and will be an important variable in what follows. We interpret $\mathcal{E}(b,z)$ as the excess elasticity households face with respect to H (which is $(1-\sigma)$), priced at the individual bundle. Since $p(b,z)$ and $\overline{\epsilon}(b,z)$ are increasing in income, $\mathcal{E}(b,z)$ is also increasing. Combining (12), (13) and (14), we obtain the following Euler equation:

$$u_c(c(b,z)) = \beta \mathbb{E}_t \left[(1+r) \frac{\mathcal{E}(b,z)}{\mathcal{E}'(b',z')} u_c(c(b',z')) \right] + \mu(b,z)\mathcal{E}(b,z), \tag{15}$$

or equivalently

$$u_c(c(b,z)) = \beta \mathbb{E}_t \left[\underbrace{\frac{1+r}{1+\pi(b',z')} \frac{\overline{\epsilon}(b,z)}{\overline{\epsilon}(b',z')}}_{\text{Effective Real Rate}} u_c(c(b',z')) \right] + \underbrace{\mu(b,z)\mathcal{E}(b,z)}_{\text{Financial Frictions}}. \tag{16}$$

Eq. (16) is the Euler equation of households of type (b,z) determining the consumption-savings decision. In the case of NH preferences, we observe that the *Effective Real Rate* depends on individual variables. In particular, it depends on the expected *individual* inflation defined as $1 + \pi(b',z') = \frac{p(b',z')}{p(b,z)}(1+$

π), and the growth of the average elasticity, so the effects of inflation are now heterogeneous across the distribution of households. Since expenditure shares are time-varying, households switch their consumption between goods with different inflation rates. That implies that whenever a household moves to a bundle with higher expected inflation, current consumption rises by more (with $u_c[c(b,z)] < 0$). Also, notice that consumption depends on the change in income elasticity out of the steady state. If the bundle becomes more elastic (i.e. if $\frac{\bar{\epsilon}(b,z)}{\bar{\epsilon}(b',z')} < 1$), there is an amplifying effect of shocks. That happens because households reallocate expenditures towards more income-elastic goods and are willing to consume fewer units on impact. However, what matters is the growth of the elasticity between today and tomorrow. Therefore, whenever this average elasticity is expected to grow over time, this implies stronger expenditure reallocation tomorrow than today to smooth consumption. Consequently, households prefer to start consuming more today. Thus, non-homoteticities contribute to consumption fluctuations through this channel.

On top of the previous channels, NH preferences affect consumption through a financial channel. For some agents, $\mu(b',z')$ is positive, inducing precautionary motives. However, the shadow price of the financial constraint now interacts with the term $p(b,z)\bar{\epsilon}(b,z)$, capturing the effective cost of living of household (b,z). Depending on the value of the latter variables, the magnitude of precautionary motives can be dampened or amplified depending on income level. That would affect the marginal propensities to consume across the income distribution. The intuition of this result is the following. Households now internalize that whenever they change the composition of their consumption bundle, they change their cost of living, $p(b,z)$. This fact also affects the cost of being close to the borrowing constraint given by $\mu(b,z)$ because reallocating to cheaper or less income-elastic bundles implies a lower need to take precautionary savings to keep high utility levels.

The latter fact directly affects the transmission of shocks in the economy: affecting precautionary motives implies distorting the curvature of consumption functions and hence affect the marginal propensities to consume of households. The final effect, though, depends on how the bundle reallocation is distributed across the population. We analyze this in the quantitative section.

Households' Distribution. The above consumption-savings problem generates a distribution of households in the space $\mathcal{B} \times \mathcal{Z}$, where $\mathcal{B}$ is determined by the borrowing constraint and $\mathcal{Z}$ by the stochastic process governing z. We denote the distribution of households by $\Psi(b,z) = \Psi(i)$, which satisfies $\int \int \Psi(i) di = 1$. Due to the recursive formulation of the problem, and given the policy functions of households, there is an operator F that maps $\Psi(i)$ onto $\Psi'(i)$

$$\Psi'(i) = F\left(\Psi(i)\right).$$

The mapping $F(\cdot)$ is a key object of our analysis because it keeps track of the distribution of households both in the steady state and along the transition path of the economy.

3.2 Workers' Union

We assume that for each task g, there is a union that decides wages W_t^g and the labor supplied n_t^g. In this setting, unions have market power as workers' tasks are in monopolistic competition. The union aggregates individual labor such that $n_t^g = \int\int n_{it}^g di$. Aggregate labor is combined in a Dixit-Stiglitz fashion, with an elasticity of demand ε_w denoting the level of market power.

We also assume nominal wage rigidities *à* la Rotemberg, in which the cost of changing wages is a convex function $\Gamma(W_{t-1}^g, W_t^g)$ measured in utility units. The optimization problem of the union is to maximize average welfare of households and, after symmetry, derives the following New-Keynesian Wage Phillips Curve (NKWPC)

$$
\begin{aligned}
(\pi_{wt}+1)\pi_{wt} - &\beta\theta_w(\pi_{wt+1}+1)\pi_{wt+1} = \\
&\frac{\varepsilon_w}{\theta_w} n_t \int\int \left\{ v'(n_t) - \frac{\varepsilon_w - 1}{\varepsilon_w} U'(c_t(b,z)) \frac{(1-\tau_t^w)W_t}{P_{it}} \frac{1-\sigma}{\bar{\epsilon}_{it}} \right\} \Psi_t(i) di,
\end{aligned}
\tag{17}
$$

where $1+\pi_{wt} = \frac{w_t}{w_{t-1}}(1+\pi_t)$ denotes wage inflation and θ_w is a parameter measuring the degree of wage rigidities. Note that, because the problem of the union depends on households' preferences, Eq. (17) depends on terms governing NH preferences.

3.3 Firms

We assume there are J sectors composed of a final good producer and intermediate producers. Final goods are a composite of a measure one of intermediates that operate in monopolistic competition, operating a CES technology with a common elasticity of substitution ε_p. Intermediate producers in sector j operate a Cobb-Douglas technology that depends on labor. The intensity of labor usage is sector-specific.

As in the case of unions, intermediate producers are subject to nominal price rigidities *à* la Rotemberg, in which the cost of changing prices is a convex function $\Theta(P_{jt-1}, P_{jt}, \theta_{jp})$ measured in units of total sectoral output, and where θ_{jp} is the sector-specific degree of price rigidities. Assuming that firms discount profits at the real interest rate and after symmetry, the optimization problem derives the following sectoral New-Keynesian Phillips Curve (NKPC)

$$
(\pi_{jt} - \bar{\pi}_j)\pi_{jt} = \frac{\varepsilon_p}{\theta_j^p}\left(\frac{mc_{jt}}{p_{jt}} - \frac{\varepsilon_p - 1}{\varepsilon_p}\right) + \mathbb{E}_t\left[\left(\frac{1}{1+r_t}\right)(\pi_{jt+1} - \bar{\pi}_j)\pi_{jt+1}\frac{p_{jt+1}y_{jt+1}}{p_{jt}y_{jt}}\right], \tag{18}
$$

where $1+\pi_{jt} = \frac{p_{jt}}{p_{jt-1}}(1+\pi_t)$ and mc_{jt} denotes marginal costs for sector j.

3.4 Monetary Authority

The monetary authority sets the nominal interest rate i_t according to a Taylor rule

$$i_t = i^* + \phi_\pi(\pi_t - \overline{\pi})$$

where we denote by ϕ_π the preference parameter for inflation.

3.5 Fiscal Policy

The government raises proportional labor income taxes τ_t^w and gives lump-sum transfers T_{it} to households. The government also issues debt denoted by B_t that has as counterpart the savings by households in liquid assets. Government debt pays the real rate. The budget constraint is then given by

$$B_{t+1} = T_t - \tau_t^w w_t N_t + (1 + r_t)B_t.$$

Additionally, we allow the government to finance the transfers with labor income taxes or debt by allowing it to delay the raise in taxes to finance the transfer today. To do so, we assume a tax smoothing parameter ρ_T governs the debt accumulation process after a fiscal transfer. Therefore, in response to a transfer, debt follows

$$dB_t = \rho_T(dT_t + dB_{t-1}).$$

We use this specification to account for debt accumulation in response to transitory government spending, as observed in the data. As we will see next, since we have a model with financial frictions and the Ricardian Equivalence does not hold, how governments finance expenditures matters for the aggregate effects of fiscal expansions.

3.6 Aggregation

Total consumption expenditures of a workers' group are given by the weighted average of their constrained and unconstrained consumers

$$C_t = \int\int p_{it} c_{it} \Psi(i) di. \tag{19}$$

We define the aggregate price index (CPI), which is the price of aggregate consumption, as $P_t = \left(\sum_j \omega_j P_{jt}^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$. We use this expression to normalize prices, and the inflation process associated with this index is the one targeted by the central bank in the case it takes spending shares constant. Therefore, aggregate inflation is defined as $1 + \pi_t = P_t/P_{t-1}$. Note also that such an index would be the one that prevails in a benchmark economy with homothetic preferences.

Then, goods market clearing implies

$$GDP_t = \sum_{j=1}^{J} P_{jt}Y_{jt} = C_t + \sum_{f=1}^{J} \Theta_{jt},$$

and the market for bonds closes

$$B_t = \int\int b_{it}\Psi(i)di.$$

4 The Role of Non-homothetic Preferences

Before turning to the quantitative application of our model, we explore the theoretical insights of NH preferences. We first explore consumption dynamics at the household level, and then we solve for an aggregate consumption (and output) equation. For this section, we assume that sectoral prices do not change over time, implying an extreme degree of price rigidities.

Consumption Dynamics at the Household Level. Starting from Eq. (8), which describes the demand of household i for good j given sectoral prices, the response of sectoral consumption to an increase in income dy_{it} is given by

$$dc_{ijt} = \underbrace{(\epsilon_j + \sigma)\omega_j \left(\frac{p_{jt}}{p_{it}}\right)^{-\sigma} (c_{it})^{\epsilon_j+\sigma-1} M_{it}^C dy_{it}}_{\text{Change through } c_{it}} + \underbrace{\sigma \left(\frac{p_{it}}{p_{jt}}\right)^{\sigma-1} \frac{1}{p_{jt}} \frac{\partial p_{it}}{\partial c_{it}} M_{it}^C c_{it}^{\epsilon_j+\sigma} dy_{it}}_{\text{Change through } p_{it}} \tag{20}$$

It can be shown that

$$dc_{ijt} = s_{ijt} \left[\frac{\sigma\bar{\epsilon}_{it} + (1-\sigma)\epsilon_j}{1-\sigma}\right] p_{it}M_{it}^C dy_{it} \tag{21}$$

where M_{it}^C denotes the marginal propensity to consume (MPC) of the household and dy_{it} the magnitude of the income shock (both taken as given for this exercise). To fix ideas, we focus on the case in which the shock comes only from fiscal transfers, so the marginal propensity to consume captures the marginal response of consumption to an infinitesimal change in transfers.

From Eq. (20), note that the response of consumption depends on good-specific characteristics (such as the price and the parameters governing preferences) and household-specific characteristics (such as the MPC or the level of aggregate consumption) through two sources: the change in consumption (as in an H economy) and the change in the individual price index (exclusive of NH). Moreover, the consumption response in the NH model differs from the H counterpart of Eq. (20) as the latter does not depend on the level of consumption and the prices index of the household. Therefore, in the H case, it does not matter the level of consumption/income of the household to determine how its sectoral consumption responds.[12]

From MPC to MPE and a Consumption Decomposition. Notice that in our model, the response of expenditures to changes in income, which is what we observe in the data, is given by $M_{it}^E = \frac{\partial E_{it}}{\partial y_t}$ (assuming $dy_{it} = dy_t \ \forall \ i$). The term M_{it}^E, which we call the *marginal propensities to spend* (MPE), is different to households' *marginal propensities to consume*, (MPC) $M_{it}^C = \frac{\partial C_{it}}{\partial y_t}$. This latter term is emphasized by Auclert et al. (2018) and is present in all models with incomplete markets and uninsurable income risk.

We can write the MPC and MPE in a Sequence-Space (Auclert et al., 2021) form to express these as *intertemporal MPC*, iMPC. Assume time is finite and given by $s \in [t, T+t-1]$ with $T < \infty$ and there are one-time, unexpected, and perfect foresight shocks (MIT shocks). The responses of consumption and expenditures can be stacked into $T \times T$ matrices, which Auclert et al. (2021) call the partial equilibrium Jacobians of households. Notice that from the individual problem of the household, we can extract the derivatives of all the individual variables to income (or fiscal transfers), namely E_{it}, c_{it}, $\bar{\epsilon}_{it}$, c_{jit}, and so on. The Jacobians for expenditures and consumption at the individual level are given by the matrices $\boldsymbol{M}_i^E$ and $\boldsymbol{M}_i^C$ which contain in rows the responses at time k of changes in income in time h given by the columns. Hence, for a path of changes in income given by $\boldsymbol{dy}$ which is a $T \times 1$ column vector, the sequence of changes in aggregate spending is given by $\boldsymbol{dE} = \int \boldsymbol{M}_i^E \boldsymbol{dy} di$. For now, we focus on a case in which the change in income only occurs in the first period t. Hence, the vector $\boldsymbol{dy}$ is zero in all elements except the first one and denote it by dy_t. Consequently, we denote the response of the endogenous variables in t as M_{it}^E for expenditures and M_{it}^C for consumption.[13]

Note that, unlike the H case (where $M_{it}^E = p_t M_{it}^C$), in the NH economy there is a gap between the MPC and the MPE since $E_{it} = p_{it}(c_{it})c_{it}$, and the price index p_{it} is household–and consumption–dependent. However, there is a mapping between M_{it}^C and M_{it}^E that helps us study the effect of having NH preferences.[14] In particular, it can be shown that

$$M_{it}^E = \frac{p_{it}\bar{\epsilon}_{it}}{1-\sigma} M_{it}^C = \mathcal{E}_{it} M_{it}^C.$$

Note that $\mathcal{E}_{it}$ is just a linear transformation of $\bar{\epsilon}_{it}$. Hence, we refer to both concepts equivalently.[15]

[12]The homothetic counterpart of Eq. (20) can be written as $dc_{ijt} = \omega_j \left(\frac{p_{jt}}{p_{it}}\right)^{-\sigma} M_{it}^C dy_t$.

[13]This is, we analyze only the first row, in time t, from matrices $\boldsymbol{M}_i^E$ and $\boldsymbol{M}_i^C$ and hence, multiplying M_{it}^E with dy_t delivers a scalar.

[14]A similar reasoning can be found in Laibson et al. (2022) considering durable and nondurable consumption.

[15]In the H model $\mathcal{E}_{it} = 1$ for all i. This can be obtained by assuming $\epsilon_j = 1 - \sigma$ and noting that $p_{it} = 1 \ \forall \ i$.

With the mapping between MPC and MPE, we can write the change in aggregate expenditures as a function of M^C_{it} and the average household elasticities

$$dE_t = \int M^E_{it} dy_t di = \int \mathcal{E}_{it} M^C_{it} dy_t \Psi(i) di. \tag{22}$$

To gain intuition, note that Eq. (22) can be written as

$$dE_t = \overline{\mathcal{E}}_t \overline{M}^C_t dy_t + cov(\mathcal{E}_{it}, M^C_{it}) dy_t, \tag{23}$$

where $\overline{\mathcal{E}}_t$ $\overline{M}^C_t$ are the row t of the aggregate Jacobian matrices of the individual $\mathcal{E}_{it}$ and consumption (in the Sequence-Space introduced by Auclert et al. (2021)).

Eq. (23) is a crucial expression determining the response of aggregate consumption to a transfer shock. It allows us to study the effects of NH preferences with respect to the benchmark H case. From our simplifying assumptions, the change in expenditures with NH preferences is characterized by two terms. On the one hand, there is an *average* effect characterized by $\overline{\mathcal{E}}_t \overline{M}^C_t$. After a positive income shock, there is a reallocation of expenditures towards more income-elastic goods in the cross-section of households. This implies an increase in the expenditure share of goods associated with a higher value of ϵ_j and, therefore, we expect $\overline{\mathcal{E}}_t > 1$ to hold. Combined with the fact that, on average, the marginal propensity to consume is positive ($\overline{M}^C_t > 0$), the average effect generates an amplification of transfer shocks.

On the other hand, there is a *cross-sectional* effect coming from the term $cov(\mathcal{E}_{it}, M^C_{it})$. After a positive income shock, we expect that the reallocation towards more income-elastic goods affects high-income households more strongly, so $\mathcal{E}_{it}$ increases in the cross-section. However, M^C_{it} moves in the opposite direction because high-income households tend to have lower marginal propensities to consume. Therefore, we expect that $cov(\mathcal{E}_{it}, M^C_{it}) < 0$ holds. We interpret this latter effect as an insurance mechanism for households. Instead of just changing their level of consumption after income shocks, households can rebalance their expenditures to keep utility as stable as possible. That interpretation of the covariance is static. A dynamic interpretation is that, in the business cycle, the distributions of $\mathcal{E}_{it}$ and M^C_{it} fluctuate, generating changes in the covariance term. In response to a transfer shock, we expect the elasticity of low-income households to rise by more than that of high-income households, lowering $cov(\mathcal{E}_{it}, M^C_{it})$ in absolute value. That may lead to an additional amplifying effect absent from models with H preferences.

In summary, the response of aggregate consumption depends on two effects, the average and cross-sectional responses. The final response of expenditures and consumption to an income shock is a quantitative result and depends on the shock's relative strength, distribution, and persistence. How does this compare with the H economy? Note that in such case, $\mathcal{E}_t = 1$ holds, and the counterpart of Eq. (23)

is just $dE_t = \overline{M}_t^C dy_t$. Therefore, in the H case, the economy has an unambiguous positive response to income shocks that only depend on the average marginal propensity to consume.

Intertemporal Responses. Now, instead of analyzing the response of variables only in period t, we turn to analyzing the sequence of responses of household variables. Auclert et al. (2018) dubbed these responses as *intertemporal MPC* for the case of consumption. These are given by the matrix multiplication (and aggregation) $\boldsymbol{dC} = \int \boldsymbol{M}_i^C dy_t di$ with $dy_t = [1, 0, ..., 0]'$. Thus $\boldsymbol{dC}$ is the sequence of responses of aggregate consumption to a one-time unitary change in income in the first period, iMPC; the equivalent object can be calculated for all the other variables, in particular, we can obtain an iMPE when analyzing expenditures. In what follows, we analyze the corresponding intertemporal responses for all household variables in the H and the NH cases to illustrate the differences between both models dynamically, in "partial equilibrium".[16] Figure 5 illustrates the intertemporal responses of household variables comparing the H and NH cases in our model calibrated for Chile and described in the next section. We show the responses of three goods (food, manufactures, and services) and the aggregate average elasticity, $\overline{\epsilon}_t$. While there are no changes in expenditure shares in the H model, we observe a sizeable effect in the NH model. There is a substantial decrease in the share of food, while a rise in manufacturing and services. Consequently, the average expenditure elasticity strongly responds to the shock in the NH model while remaining constant in the H case. Therefore, Figure 5 shows that our model can replicate the empirical evidence presented in Section 2: the share of food in total consumption goes down, and the share of services and manufactures goes up after an income shock.

FIGURE 5: Change in Expenditure Shares and Income Elasticity to a One-Time Unitary Fiscal Transfer

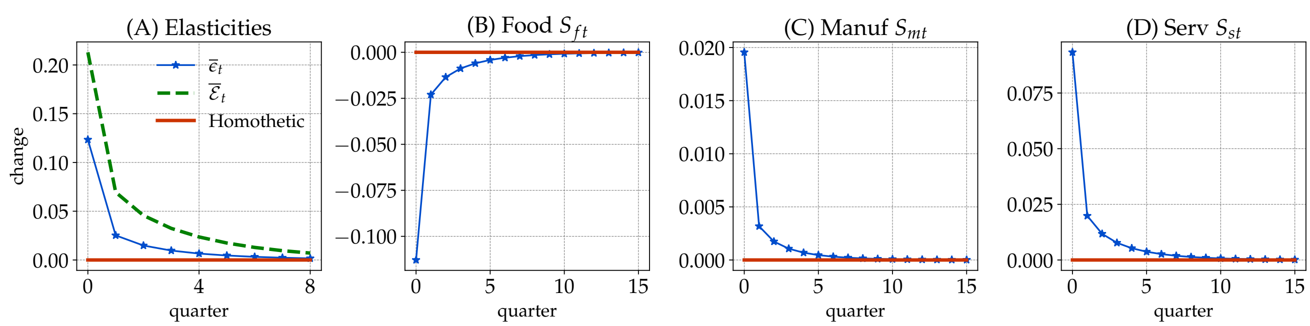

NOTES: This figure shows the dynamic response of elasticities and the share of food, manufacturing, and services to a one-time increase in household transfers. These responses are analogous to the partial Jacobians proposed by Auclert et al. (2021) for consumption they call *iMPCs*. We show the H and the NH cases.

[16] Auclert et al. (2018) show that in models with heterogeneous agents, these intertemporal MPCs arise naturally, in contrast with other types of modeling like Two-Agent models (Bilbiie, 2008) in which the responses to one-time fiscal transfers have only effects in the period of the shock.

The responses in expenditure shares and the average elasticity affect the responses of consumption in the different goods and the response of expenditures, the iMPE in the NH case. Figure 6 shows that, as we expected, the responses of the different goods follow the income elasticities and the response of the expenditure shares. In the NH case, food consumption rises by less than the rise in manufactured goods and services in our calibration. Notice also that in the NH case, the iMPC is slightly lower than the iMPC of the H case, but the iMPE lies above the iMPC in H. This result suggests that the response of the elasticity shown in Panel (A) of Figure 5 contributes to the rise in expenditures, hence, generating amplification.

FIGURE 6: Changes in Expenditures and Consumption to a One-Time Unitary Transfer (iMPC, iMPE)

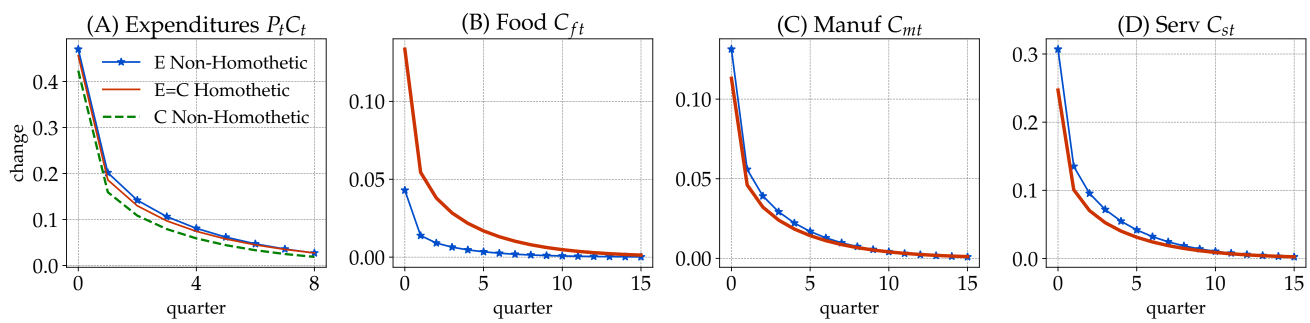

NOTES: This figure shows the dynamic response of consumption, expenditures, and consumption of the different goods to a one-time increase in household transfers. These responses are analogous to the partial Jacobians proposed by Auclert et al. (2021) for consumption they call *iMPC*s. We show the H and the NH cases.

The Effect of NH on the Keynesian Cross. While Figure 6 already shows that our mechanism–for the calibration assumed in this paper– generates amplification of the effects of transfers through higher iMPE, this still is a general equilibrium question. To make progress, we write Eq. (23) assuming $dE_t = dY_t$ and $dy_t = dY_t + dT_t$ as

$$dY_t = \overline{\mathcal{E}}_t \overline{M}_t (dY_t + dT_t) + cov(\mathcal{E}_{it}, M_{it})(dY_t + dT_t)$$

where we can write a general equilibrium counterpart in the form of a "Keynesian Cross" (Auclert et al., 2018) as[17]

$$dY_t = \frac{\overline{\mathcal{E}}_t \overline{M}_t + cov(\mathcal{E}_{it}, M_{it})}{1 - \overline{\mathcal{E}}_t \overline{M}_t - cov(\mathcal{E}_{it}, M_{it})} dT_t. \tag{24}$$

[17]This can be obtained by noting that $dy_t = d[(1 - \tau_t^w) w_t N_t] + dD_t + dT_t = dY_t + dT_t$ assuming that the fiscal transfer is fully debt-financed and $\tau^w d[w_t N_t]$ goes to debt payments as well. We relax these assumptions in the quantitative section.

Thus, NH preferences contributes to amplification by having a combination of a large $\overline{\mathcal{E}}_t$ and a low $cov(\mathcal{E}_{it}, M_{it})$, but also due to intertemporal changes in these two terms: there is stronger amplification whenever there are large changes in these components.

In the next section, we study the transmission of transfer shocks in the full-blown model and show on what the effects of NH depend.

5 The Transmission of Fiscal Shocks

This section studies the aggregate implications of non-homothetic preferences after an income shock. We start by describing the calibration of the model to then present the responses of the economy under different counterfactual scenarios.

5.1 Calibration

Estimation of Income Elasticities. As noticed by Comin et al. (2021), the model's predictions for observable variables remain invariant to any scaling of all income elasticities and taste shifters (ϵ_j and ω_j) by a constant factor. Therefore, we can normalize all these parameters relative to a base good. Let $j = v$ denote such base good, which will be normalized to one (i.e., $\epsilon_v = \omega_v = 1$). This implies that we can write the real consumption index as $c_{it} = s_{ivt}(p_{jt}/E_{it})^{\sigma-1}$. Substituting this expression back in Eq. (9) for any $j \neq v$, the expenditure share in good j relative to the base good by household i in period t can be written as:

$$\log\left(\frac{s_{ijt}}{s_{ivt}}\right) =(\epsilon_j - 1)\log(s_{vt}) + (1-\sigma)\log\left(\frac{p_{jt}}{p_{vt}}\right) + (\epsilon_j - 1)(1-\sigma)\log\left(\frac{E_{it}}{p_{vt}}\right), \tag{25}$$

for any $j \neq v$, which defines a $J-1$ system of demand equations. The key element to notice from Eq. (25) is that it provides an expression for the consumption shares of all other goods in terms of observable variables. For our empirical estimation, we set manufactures as the baseline good ($v = m$) and normalized those values to one ($\epsilon_m = \omega_m = 1$).

We estimate this empirical specification with the cross-sectional data presented in Section 2 to replicate the expenditure patterns observed across the income distribution. As Comin et al. (2021), we assign a percentile-specific price for the three goods, which imperfectly captures the fact that different households might not face the same prices. We impose that the parameters are the same across the $J-1$ equations. The estimation is carried out using Feasible Generalized Nonlinear Least Squares (FGNLS) as Herrendorf et al. (2013) and Cravino and Sotelo (2019).[18] As in Comin et al. (2021) and Cravino and Sotelo (2019), the

[18]For the estimation, we constraint the elasticities to be positive to ensure that the consumption aggregator is concave.

identification assumption is that shocks to income and relative prices are not correlated to changes in demand shifters, ω_j, so preferences do not change over time, other than by the income effect.

Table 3 reports the results. The elasticity of substitution σ is significantly below one and close to zero, implying a high degree of complementarities in consumption. On the other hand, the income elasticity of food is close to zero, while the income elasticity of services is above one. These results indicate that services are more income elastic than manufactures and food.[19]

TABLE 3: Demand system estimates

	Coefficient	Std. Error
σ	0.271***	(0.023)
ϵ_f	0.000	$(\cdot)$
ϵ_s	1.113***	(0.036)
Observations	100	

NOTES: This table presents the estimates of the demand system given by Eq. (25). σ denotes the elasticity of substitution between goods, while ϵ_j denotes the income elasticity of good $j \in \{f, s\}$. Robust standard errors reported in parenthesis. *, ** and *** denote statistical significance at the 1, 5 and 10% levels, respectively.

Households. We set the elasticity of intertemporal substitution equal to one ($\gamma = 1$). We calibrate the disutility of labor ψ to match a normalized level of labor given by $N = 1$ in steady state, and the Frisch elasticity of labor is equal to one ($\varphi = 1$). We calibrate the discount factor and the average bond holdings to match an interest rate of 0.5% quarterly and a share of hand-to-mouth workers of 0.39 (García et al., 2023). This calibration strategy delivers total liquid holdings of 16.5 percent of GDP and a discount factor of $\beta = 0.94$. Total liquid holdings are somewhat lower than García et al. (2023) get but similar. Finally, we assume an ad-hoc rule of profits distribution in which we distribute profits proportional to the idiosyncratic productivity as in Kaplan et al. (2018).

We calibrate the income risk process to match the cross-sectional variance of the first difference of log income (0.19) and the variance of the log of income (0.72) at a quarterly frequency.[20] We discretize this process in eleven points using the Rouwenhorst method. To estimate the idiosyncratic income process, we use administrative data with information on the universe of formal workers at a quarterly frequency. Finally, we assume workers' unions face an elasticity of substitution equal to $\varepsilon = 10$ and calibrate the adjustment cost to obtain a slope of the New Keynesian Wage Phillips curve equal to 0.1.

[19] For comparison purposes, note that Comin et al. (2021) finds $\sigma = 0.26$, $\epsilon_f = 0.2$ and $\epsilon_s = 1.65$ for the U.S., using panel data from the Consumption Expenditure Survey (CEX) for the period 1999-2010.

[20] Also reported in García et al. (2023).

Firms. For consistency with our empirical results, we assume three sectors: food, manufactures, and services. For simplicity, we assume all firms face the same elasticity of demand, ε_j, which equals 10. In the baseline calibration, we set the adjustment cost parameters to obtain slopes of the Phillips curves equal to 0.1 in all goods. Additionally, we assume decreasing returns to scale in labor, setting $\alpha = 0.33$.

To isolate the effects from consumption to output, we assume firms are equal in how they produce but not in the demands they face. The last parameters we have left to calibrate this economy are the taste parameters ω's, which we calibrate to match the average expenditure shares on each group of goods in steady state. From Table A.1 we target an expenditure share in food, manufactures, and services of 0.204, 0.185, and 0.611. Our calibrated parameters are $\omega_f = 0.274$, $\omega_m = 0.228$ and $\omega_s = 0.553$.

Monetary and Fiscal Policy. The Taylor rule only responds to inflation. In both the homothetic and non-homothetic cases, it targets the average inflation given by the aggregate price index. We set $\phi_\pi = 1.5$.

For fiscal policy, we assume that the government finances transfers and interest rate payments with distortionary labor income taxes. Those transfers are lump-sum and equal to all workers so we abstract from progressivity and we assume the persistence of the transfers to be 0.75, which means a half-life of two quarters. We calibrate the level of the proportional labor tax (τ^w) to satisfy the government's budget constraint to a value of $\tau^w = 0.145$. We target a value of aggregate fiscal transfers equal to ten percent of GDP in steady-state. Finally, we set the autoregressive process governing the dynamics of fiscal debt to $\rho_T = 0.75$, which implies that the government accumulates debt initially and starts decreasing debt in the following periods by raising labor income taxes.

Solution Method. The steady-state equilibrium is obtained by solving the model equilibrium for seven parameters and prices. To solve the value function, we use Carroll (2005) endogenous grid method, a fast and accurate algorithm to solve these problems. Finally, we use a Newton method to solve for β, B, φ, p_2, p_3, ω_1, and ω_2 which satisfy equilibrium in steady-state. To solve for the dynamics of the model with aggregate shocks, we follow Auclert et al. (2021), who propose to write the model in its Sequence-Space form and linearize it around this system of equations. The method relies on the fact that any model without aggregate uncertainty can be written as a sequence of equations. This is, if we assume shocks are unexpected and the path is known (MIT shocks), we can write the system as a sequence of equations in the transitional dynamics. This system of equations given by $T \times M$, with T the horizon of the transition and M the number of equations to solve, can be linearized around the steady state. This linearization leads to jacobians of all variables with respect to others, and a composition of these jacobians delivers the impulse-response functions. This method is based on the result by Boppart et al. (2018) that shows that the solution from transitional dynamics is close to the linearized solution if shocks are small. The method is fast, accurate, and robust relative to methods like Reiter (2009). We refer the reader to the paper for more details on the method.

5.2 Aggregate and Sectoral Responses to Fiscal Transfer Shocks

This section presents the sectoral and aggregate responses to a fiscal transfer shock under our baseline calibration. The experiment is as follows. We let the government increase the aggregate lump sum transfer and distribute it evenly among households (i.e., all households receive the same amount). The rise in transfers is one percent of GDP on impact. In all our analyzes, we compare the response of the non-homothetic (NH) model with its homothetic (H) counterpart, where the H counterpart follows by setting $\epsilon_j = 1 - \sigma \ \forall j$.

Figure 7 presents the economy's responses after the fiscal transfer shock. Panels (A)-(C) show the aggregate response to sectoral goods, while panel (D) shows the response of GDP after the shock. The blue-star lines correspond to the responses in the NH model, while the red-solid lines are the responses in the H counterpart.

Since we study a shock that increases the income level in the economy, all the sectoral and aggregate responses increase the level of consumption and activity, regardless of the distinction between an NH or an H model. However, there are important quantitative differences between the two economies. Regarding sectoral consumption (panels A-C), we observe a less pronounced response in food expenditures in the NH model but a stronger response in manufactures and services expenditures relative to the H case. This is in line with our previous discussion, in which after a positive income shock, there is a rebalancing in expenditures from less to more income-elastic goods. Consistent with our empirical evidence in Section (2), food is less income-elastic than manufactures, which in turn is less income-elastic than services. These differences have important implications for aggregate activity. In panel (D), we observe that the response of GDP is about 60 percent larger and more persistent in the NH economy than in its H counterpart. Such difference can also be noticed by looking at panel (E), which shows the differences in the response between GDP in the NH case relative to the H counterpart.

FIGURE 7: IRFs under Baseline Calibration

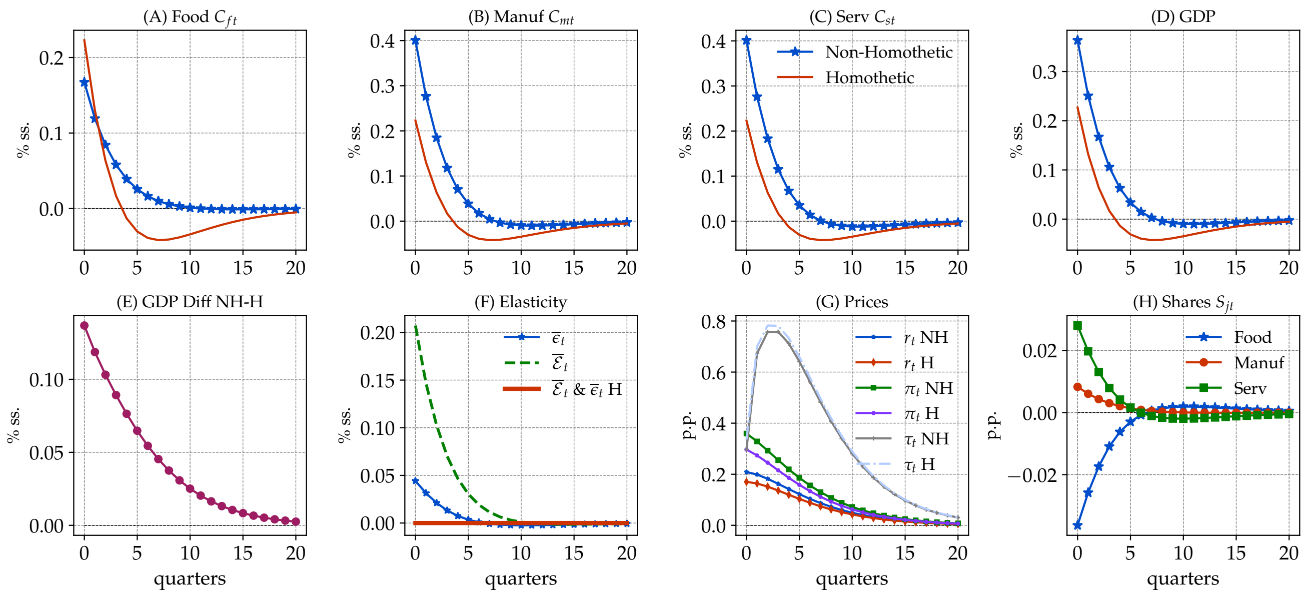

NOTES: This figure shows the sectoral and aggregate responses to a fiscal transfer shock under our baseline calibration.

Why is this the case? The key feature of the NH model comes from the dependency of expenditure shares on the income level. Therefore, we expect that those expenditure shares react in a meaningful way after an income shock. Panel (H) in Figure 7 shows that this is the case, with a positive response in expenditure shares in manufactures and services (larger response in the latter) and a negative response in expenditure shares of food in the NH model. Also, even though relative prices are also changing in general equilibrium, the response of expenditure shares is muted in the H version of the model due to our assumption of an equal degree of price rigidities across sectors.

The response of sectoral expenditure shares has implications over the average income-elasticity at the household level, $\overline{\epsilon}_{it} = \sum_j s_{ijt}\epsilon_j$, which determines the final impact of the real interest rate determining consumption, as in Eq. (16). As shown in panel (F) in Figure 7, the average income-elasticity strongly responds to the shock, given the reallocation towards more income-elastic goods. This is not the case in the H model, in which, by definition, such elasticity is constant. Therefore, we expect an increase in the cross-sectional average income elasticity captured by the $\overline{\mathcal{E}}_t$, which amplifies the response to shocks due to the average marginal propensity to consume.

The differences between the NH model and its H counterpart are reflected in the response of real variables and the response of nominal ones. As shown in panel (G) of Figure 7, even though the response of fiscal variables is similar in both models, there are sizeable differences in the response of aggregate inflation. In the NH model, the response of inflation is stronger and more persistent than in the H model, which translates into a stronger response of monetary policy to contain inflation. Therefore, even with a more countercyclical monetary policy, we find that NH preferences amplify the response of GDP to fiscal transfers.

Our analysis does not rely on sectoral-specific shocks to generate asymmetric sectoral responses, as we study an aggregate shock. Therefore, our results shed light on the importance of analyzing the role of income in understanding changes in demand composition over the business cycle and how such changes account for the response and persistence of aggregate variables. Hence, we conclude that considering a mechanism that generates expenditure reallocation can be important for the response of aggregate variables and the response of inflation, on top of the mechanism proposed by heterogeneous price rigidities. This point has been highlighted by Ferrante et al. (forthcoming) but considering an exogenous bundle-reallocation shock.

On the Role of Shocks' Persistence The aggregate response of the economy after an income shock crucially depends on how households reallocate expenditures. This force is intertemporally driven by the persistence of shocks and how they change the average income-elasticity, $\overline{\epsilon}_{it}$ at the household level. Also, as discussed in Section 4, the response in the (cross-sectional) average income-elasticity (also by $\overline{\mathcal{E}}_t$), has implications over the static average and cross-sectional forces, which might amplify or dampen shocks. How important is the persistence of shocks in driving aggregate responses?

Figure 8 analyzes the aggregate responses of the economy, considering a lower persistence of the transfer shock, from 0.75 in the baseline scenario to only 0.25. Even though all the qualitative insights of the baseline economy remain, we observe a higher response of GDP to a less persistent transfer. In the baseline, the government has to raise taxes more strongly because there are more periods in which the transfer occurs. Hence the transfer is less expansionary overall. This is a result of our calibration; if we assume a more debt-financed transfer this result may not hold.

However, this exercise generates a complementary result which is that due to the less persistent transfer, the elasticity remains high for a shorter time than in the baseline, and that generates a weaker contribution of NH. As Eq. (23) shows, there is an interaction between $\overline{\mathcal{E}}_t$ and the intertemporal MPCs that drive the effects of the transfers. Such interaction is weaker when the shock is less persistent, implying a smaller contribution of NH to the response in expenditures. This is observed in panel (F) in Figure 8, which shows that the elasticity has a half-life of about one period while that of the baseline calibration is about two periods. This is the main driver of the amplification observed in the baseline analysis, which we break in this exercise.[21]

[21] We tried for several calibrations of ρ_T (0.9 and 0.95) and the persistence of transfers, and we conclude the same.

FIGURE 8: IRFs with Lower Persistence in Transfers

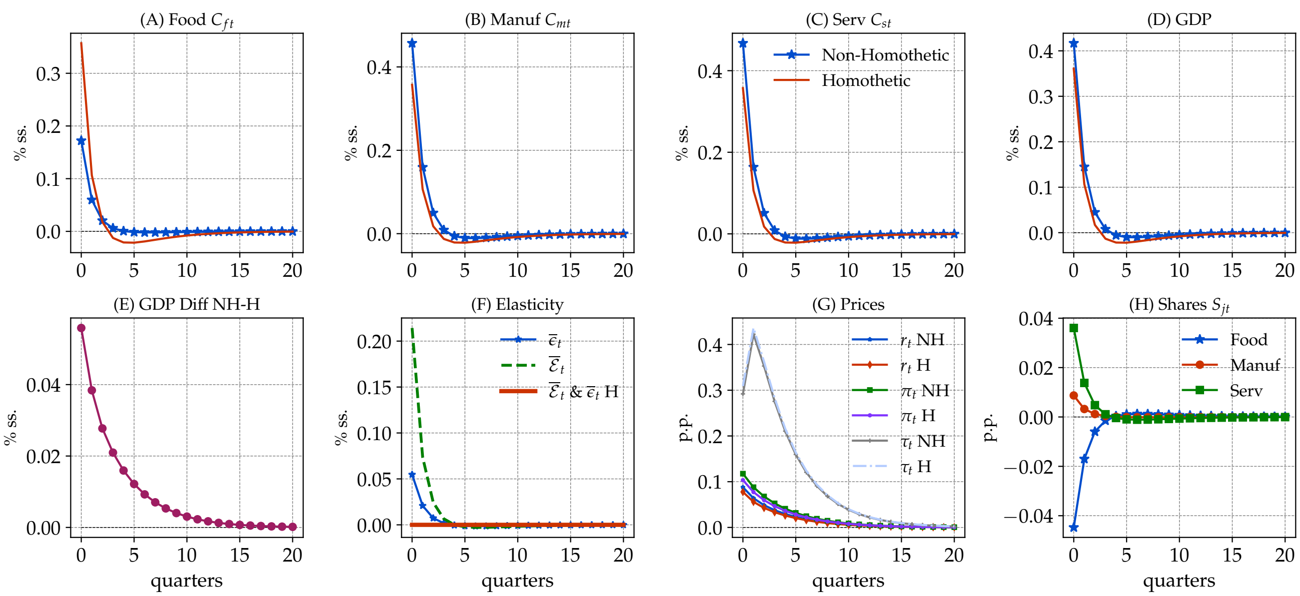

NOTES: This figure shows the sectoral and aggregate responses to a fiscal transfer shock under a less persistent one.

With Countercyclical Labor Income Inequality. Our baseline experiment assumes that labor income and transfers affect all households similarly, so the only sources of heterogeneity (over the business cycle) are bond and profit holdings. A more realistic assumption would be that low-income households face more procyclical labor income fluctuations. Hence, there is a stronger redistribution of resources in the cycle towards households with high MPCs, and with lower income elasticities.[22]

Figure 9 investigates the effects of fiscal transfers in which the labor income of low-income households responds more strongly to shocks. In this calibration, we assume that hours, instead of being the same for all households, vary with household idiosyncratic productivity. We assume the following specification for this behavior $N_{it} = N_t \exp\{z_i \xi (N_t - N_{ss})\} / N_0$,[23] with $\xi = -2$, which generates a more procyclical response of labor income in the lower part of the labor income distribution.

The results of Figure 9 show that with more countercyclical inequality, the response of GDP gets amplified in both the NH and the H cases. This is, there is amplification because income is distributed more procyclically towards households with higher MPCs. More interestingly, the amplification we observe in the baseline economy is augmented by the countercyclical inequality. This is because the response of the cross-sectional average elasticity is stronger than in the baseline. These results imply that the covariance between the elasticities and the MPCs is also becoming less negative, contributing to

[22]Evidence for this fact in Chile can be found in Aldunate et al. (2023).

[23]With $N_0 = \int \exp\{z_i \xi (N_t - N_{ss})\} dz$

the amplification effects. Therefore, we conclude there is more amplification of NH preferences through expenditures reallocation when the shock is distributed more strongly in high-MPC households.

FIGURE 9: IRFs with Countercyclical Labor Income Inequality

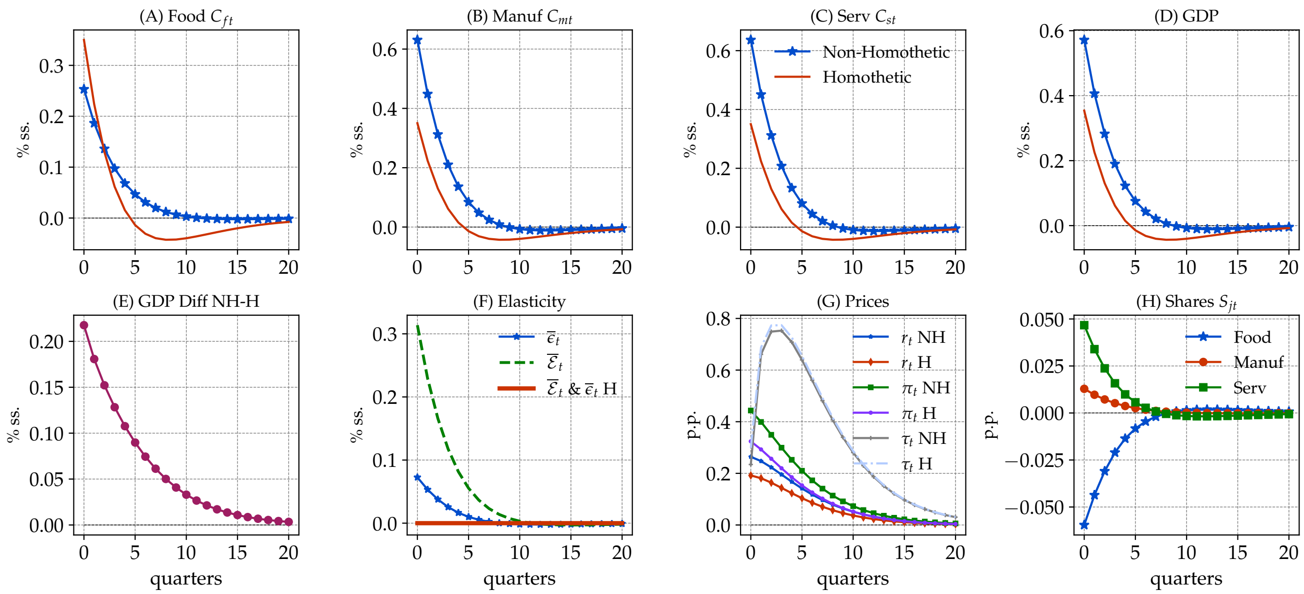

NOTES: This figure shows the sectoral and aggregate responses to a fiscal transfer shock under a countercyclical labor income inequality.

On the Role of Monetary Policy Our previous exercises were conducted by assuming a monetary policy authority following the Taylor principle and reacting more than one-to-one to inflation. This behavior is crucial to contain the demand-driven pressures in the economy by affecting the effective real rate presented in Eq. (16). What are the implications of having a looser monetary policy?

We consider the aggregate responses of the economy to the fiscal transfer shock, with a response to inflation in the Taylor rule of $\phi_\pi = 1$. Figure 10 presents the results. A looser monetary policy implies a lower real rate response, which is strongly amplified by the variation in the average income elasticity. This has a strong impact on sectoral reallocation and expenditures, and aggregate GDP, with responses that are more than two times larger relative to the baseline scenario. Because these mechanisms are not present in the H model, the differences in the response of aggregate GDP between the two models are also stronger and more persistent. Moreover, the response of aggregate inflation doubles and is highly persistent over time.

FIGURE 10: IRFs under Loose Monetary Policy

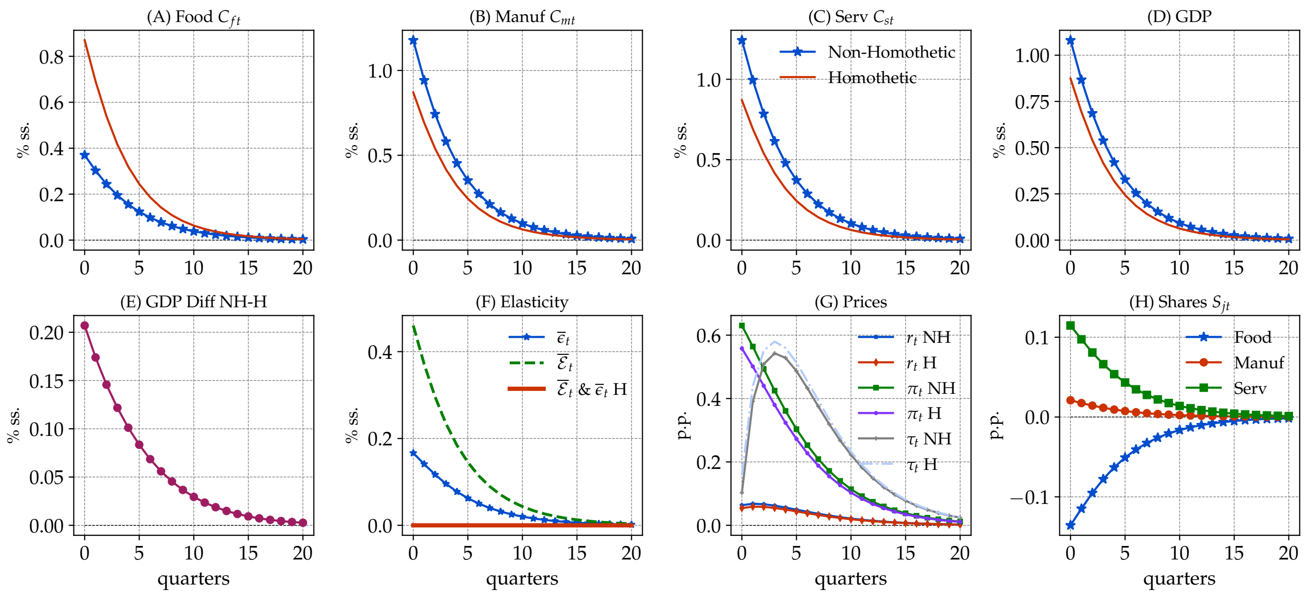

NOTES: This figure shows the sectoral and aggregate responses to a fiscal transfer shock under a looser monetary policy.

Heterogeneity in Price Rigidities. Finally, we study the effects of fiscal transfers considering an alternative explanation for the responses of expenditure shares, the fact that the different types of goods have different price rigidities. In particular, we consider the case in which services have more rigid prices than manufacturing, and manufacturing goods have more rigid prices than food. This evidence is stressed by Nakamura and Steinsson (2008). In principle, this may generate more real effects on services production, generating changes in expenditure shares more concentrated in services.

However, this interpretation might be imprecise. Recall that aggregate expenditure shares are defined as $P_{jt}C_{jt}/GDP_t$. Hence, expenditure shares vary both due to relative changes in prices and quantities, and it is not clear that more rigidities in more income-elastic goods generate more amplification or even the cyclicalities of the expenditure shares we observe in the data. Recall that in our baseline exercise, we assume the same slope of the Phillips curve for the three sectors equals 0.1. Now we break this assumption and assume that the slope of the NKPC of food is 0.15, while for manufacturing and services is 0.1 and 0.05, respectively. This alternative calibration captures the idea that services have more rigid prices than manufactures and food.

Figure 11 shows the result of this exercise. Notice that, on impact, we obtain a similar effect in the response of expenditure shares in the NH model, as shown in panel (H). However, the expenditure shares in the H case (see panel G) are inconsistent with the empirical facts presented in Section 2: in the H economy we observe that expenditures in food and manufacturing go up while expenditures in services go down, while in the data we show that the expenditure shares have the opposite patterns for food and services. Hence, we conclude that the most likely driver of the expenditure shares are income elasticities and not heterogeneity in price rigidities when we study the effects of fiscal transfers.

FIGURE 11: IRFs with Heterogeneity in Price Rigidities

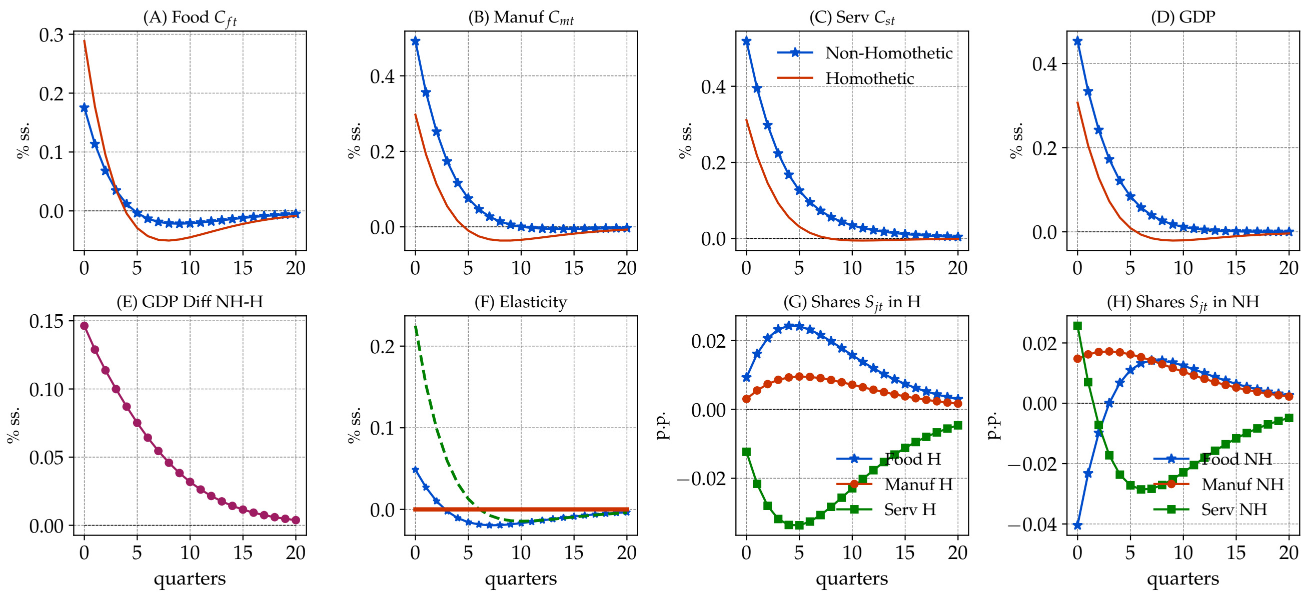

NOTES: This figure shows the sectoral and aggregate responses to a fiscal transfer shock under heterogeneous price rigidities.

6 Conclusion

In this paper, we analyze the role of income heterogeneity in determining aggregate consumption. We provide evidence on how expenditures are distributed across the income distribution, showing that richer households spend relatively more on services and less on food and manufacturing. Moreover, we show that these expenditure shares change over the business cycle and after income shocks, with stronger responses in the share of services than manufactures or food. This is, we show a reallocation towards more income-elastic goods over the business cycle.

Motivated by this evidence, we build a Heterogeneous Agent New Keynesian model considering non-homothetic preferences (NH) to study how income determines expenditure patterns and consumption. Our key result is that NH preferences affect the consumption response to fiscal transfers both intratemporally (by reallocating consumption towards more income-elastic goods) and intertemporally (through a real interest rate channel and a financial frictions channel). Furthermore, our quantitative results indicate that the conduct of monetary policy and the persistence of shocks are crucial to understanding the amplifying effects of NH preferences.

Our paper contributes to a growing literature analyzing the role of income shocks in determining business cycles and how microeconomic responses shape aggregate responses. In particular, we contribute by studying the role of income in determining expenditure shares, the reallocation of consumption baskets, and how these elements shape aggregate responses. The endogenous spending reallocation mechanism that we propose also contributes, from a theoretical perspective, to the literature on the cyclical behavior of marginal propensities to consume (see Gross et al., 2020) by incorporating an additional margin arising from a time-varying aggregate income elasticity.

References

ALDUNATE, R., A. BLANCO, A. FERNÁNDEZ, M. GIARDA, AND G. NAVARRO (2023): "The Cross-Sectional Labor Market Dynamics After a Foreign Shock," *mimeo Central Bank of Chile*.

AUCLERT, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–67.

AUCLERT, A., B. BARDOCZY, M. ROGNLIE, AND L. STRAUB (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375—-2408.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): "The Intertemporal Keynesian Cross," Working Paper 25020, National Bureau of Economic Research.

BAQAEE, D., E. FARHI, AND K. SANGANI (2021): "The Supply-Side Effects of Monetary Policy," Working Paper 28345, National Bureau of Economic Research.

BILBIIE, F. O. (2008): "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of Economic Theory*, 140, 162 – 196.

BOPPART, T., P. KRUSELL, AND K. MITMAN (2018): "Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative," *Journal of Economic Dynamics and Control*, 89, 68–92.

CARROLL, C. (2005): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economic Letters*, 91, 312–320.

CLAYTON, C., X. JARAVEL, AND A. SCHAAB (2018): "Heterogeneous Price Rigidities and Monetary Policy," *mimeo*.

COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): "Structural Change with Long-run Income and Price Effects," *Econometrica*, 89, 311–374.

CRAVINO, J., T. LAN, AND A. LEVCHENKO (2020): "Price Stickiness Along the Income Distribution and the Effects of Monetary Policy," *Journal of Monetary Economics*, 110, 19–32.

CRAVINO, J. AND A. LEVCHENKO (2017): "The Distributional Consequences of Large Devaluations," *American Economic Review*, 107, 3477–3509.

CRAVINO, J. AND S. SOTELO (2019): "Trade-Induced Structural Change and the Skill Premium," *American Economic Journal: Macroeconomics*, 11, 289–326.

FERRANTE, F., S. GRAVES, AND M. IACOVIELLO (forthcoming): "The Inflationary Effects of Sectoral Reallocation," *Journal of Monetary Economics*.

FUSTER, A., G. KAPLAN, AND B. ZAFAR (2020): "What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans," *The Review of Economic Studies*, 88, 1760–1795.

GALÍ, J. (2011): "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, 9, 436–461.

——— (2013): "Notes For a New Guide to Keynes (i): Wages, Aggregate Demand, and Employment," *Journal of the European Economic Association*, 11, 973–1003.

GARCÍA, B., M. GIARDA, AND C. LIZAMA (2023): "The Role of Progressivity on the Economic Impact of Fiscal Transfers: A HANK for Chile," *mimeo Central Bank of Chile*.

GROSS, T., M. J. NOTOWIDIGDO, AND J. WANG (2020): "The marginal propensity to consume over the business cycle," *American Economic Journal: Macroeconomics*, 12, 351–384.

GUREN, A., A. MCKAY, E. NAKAMURA, AND J. STEINSSON (2020): "What Do We Learn from Cross-Regional Empirical Estimates in Macroeconomics?" *NBER Macro Annual*, 35.

HANOCH, G. (1975): "Production and Demand Models with Direct or Indirect Implicit Additivity," *Econometrica*, 43, 395–419.

HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103, 2752–2789.

JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96, 1589–1610.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

LAIBSON, D., P. MAXTED, AND B. MOLL (2022): "A Simple Mapping from MPCs to MPXs," NBER Working Papers 29664, National Bureau of Economic Research, Inc.

MATSUYAMA, K. (2019): "Engel's Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade," *Econometrica*, 87, 497–528.

MIAN, A., K. RAO, AND A. SUFI (2013): "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*, 128, 1687–1726.

MIAN, A. AND A. SUFI (2009): "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis," *The Quarterly Journal of Economics*, 124, 1449–1496.

MISRA, K. AND P. SURICO (2014): "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs," *American Economic Journal: Macroeconomics*, 6, 84–106.

NAKAMURA, E. AND J. STEINSSON (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 123, 1415–1464.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103, 2530–53.

PASTEN, E., R. SCHOENLE, AND M. WEBER (2020): "The propagation of monetary policy shocks in a heterogeneous production economy," *Journal of Monetary Economics*, 116, 1–22.

PATTERSON, C. (2019): "The Matching Multiplier and the Amplification of Recessions," *mimeo MIT*.

REDDING, S. AND D. WEINSTEIN (2019): "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences," *The Quarterly Journal of Economics*, 135, 503–560.

REITER, M. (2009): "Solving Heterogeneous-agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 33, 649–665.

ROMERO, D. (2022): "Market Incompleteness, Consumption Heterogeneity and Commodity Price Shocks," *mimeo*.

ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

SCHAAB, A. AND S. Y. TAN (2023): "Monetary and Fiscal Policy According to HANK-IO," *mimeo Toulouse*.

A Empirical Appendix

This appendix provides additional empirical results. First, we study the robustness of our static evidence presented in Section 2.2. Then, we present suggestive evidence on the role of non-homothetic preferences across a panel of countries, which helps us to give external validation to our approach. After that, we present comparisons between our micro-level data and aggregate data on expenditure patterns, as well as the representativity of income at the municipality level. Finally, we provide additional evidence related to the time-varying behavior of expenditure shares presented in Sections 2.3 and 2.4.

A.1 Additional Static Evidence

A.1.1 Disaggregated Expenditure Patterns

Table A.1 presents additional evidence for selected percentiles in the income distribution and a more disaggregated level of consumption categories. As shown in the table, the key message of the aggregate expenditures division remains when we decompose into a more disaggregated level of consumption. This is, expenditure shares in consumption categories within food and beverages (services) groups are decreasing (increasing) across the income distribution.

A.1.2 Consumption Patterns Across Labor Income Distribution

In the main text, the cross-sectional heterogeneity is constructed considering all sources of income of households. Figure A.1 compares consumption expenditures between labor income and total income distributions. As can be seen, both present a similar picture, in which low-income households (measured either by labor income or total income) spent a larger fraction of their income in food and beverages, while richer households spent more on services. While both distributions are closely correlated for those goods (above 90 percent), larger differences are presented for manufactured goods, in which the correlation is just 23 percent.

A.1.3 Expenditure Patterns Across Different Waves of the Consumption Expenditure Survey

Our baseline evidence uses the most recent wave of the Consumption Expenditure Survey of 2017. How consistent are those patterns across time?

Figure A.2 compares the distribution of consumption expenditures across different waves of the Consumption Expenditure Survey, EPF. The baseline sample corresponds to EPF VIII of the year 2017 (x-axis on each panel), while the alternative samples corresponds to EPF V of 1996 (green squares on each panel), EPF VI of 2006 (red triangles on each panel), and EPF VII of 2014 (blue dots on each panel). On

TABLE A.1: Consumption expenditure of households

Code	Consumption division	Income percentile					Average
		P10	P25	P50	P75	P90	
Panel A: Food and beverages							
01	Food and non-alcoholic beverages	0.325	0.276	0.232	0.160	0.110	0.186
02	Alcoholic beverages, tobacco	0.021	0.020	0.016	0.021	0.016	0.018
	Total	0.346	0.295	0.248	0.181	0.126	0.204
Panel B: Manufactures							
03	Clothing and footwear	0.047	0.043	0.037	0.032	0.033	0.035
04	Housing, water, electricity, gas and other fuels	0.100	0.098	0.110	0.078	0.068	0.087
05	Furnishings, household equipment	0.043	0.037	0.042	0.060	0.086	0.062
	Total	0.190	0.178	0.189	0.170	0.187	0.185
Panel C: Services							
04.1	Rentals	0.043	0.034	0.035	0.053	0.062	0.053
06	Health	0.036	0.076	0.094	0.072	0.069	0.072
07	Transport	0.093	0.137	0.122	0.165	0.159	0.156
08	Communication	0.050	0.058	0.069	0.063	0.044	0.052
09	Recreation and culture	0.058	0.053	0.064	0.115	0.133	0.086
10	Education	0.088	0.053	0.063	0.032	0.051	0.046
11	Restaurants and hotels	0.045	0.046	0.043	0.069	0.078	0.067
12	Miscellaneous goods and services	0.052	0.070	0.073	0.080	0.090	0.078
	Total	0.464	0.527	0.563	0.649	0.687	0.611

NOTES: This table presents the share in consumption expenditures for households in selected percentiles of the income distribution, considering the 12 division of expenditure groups. Codes corresponds to the 12 divisions in the Classification of Individual Consumption by Purpose (COICOP). Panel A presents the decomposition for Food and Beverages categories. Panel B presents the decomposition for Manufactures, housing and utilities. Panel C presents the decomposition for Services. Each column denotes percentiles 10, 25, 50, 75 and 90, and average consumption, respectively.

each panel, every point shows the expenditure share in the respective kind of good in EPF VIII against other waves of the survey. As can be seen, with the exception of manufactured goods, the expenditure patterns documented in the main text are relatively stable over time. In the case of food and beverages,

FIGURE A.1: Consumption patterns across: aggregate vs labor income distribution

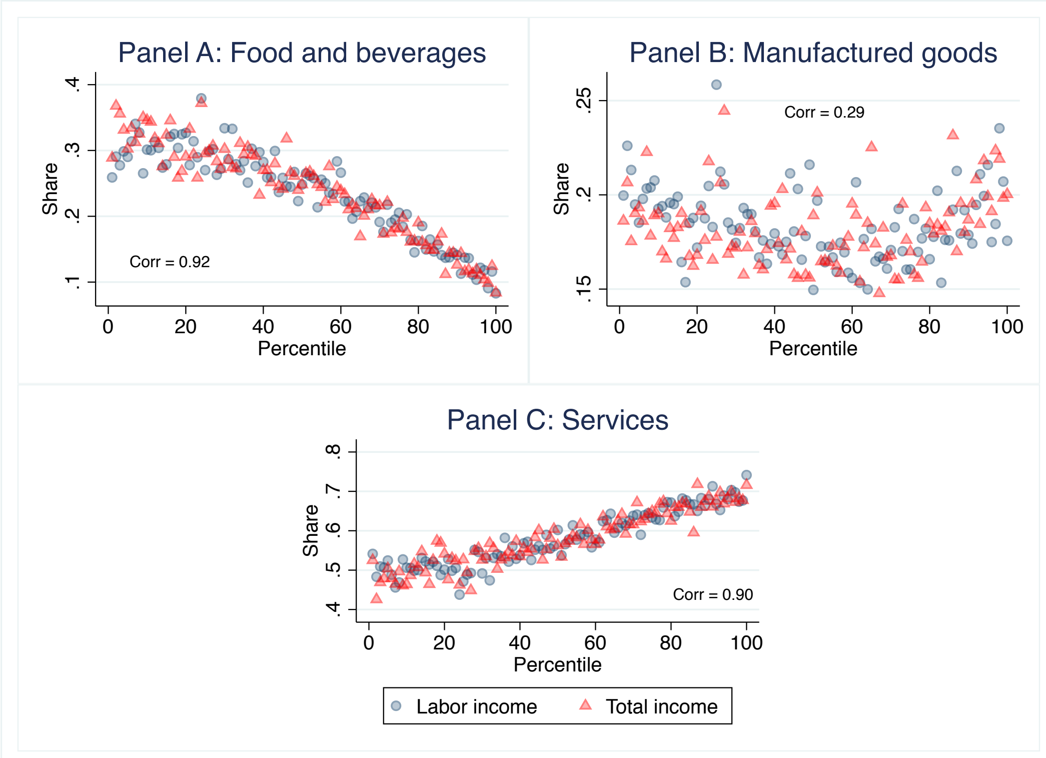

NOTES: This figure compares the share in consumption expenditures for households in each percentile of the income distribution vs percentiles in labor income distribution, considering three aggregate expenditure groups. On every panel, each dot/triangle corresponds to a percentile.

correlations with respect to the baseline year are above 0.9, while for services they are above 0.8.[24] Note that the largest differences in levels are observed with respect to EPF V, because of the distance in time with EPF VIII (20 years). In particular, food and beverages account for a greater fraction of expenditures in EPF V, while the opposite happens in services.

A.2 Expenditure Patterns in Other Countries

Our empirical evidence focuses in the case of Chile. How representative are those expenditure patterns for other countries?

While comparing expenditure patterns in the cross-section of different countries is a difficult task, we do our best to provide further evidence supporting the generality of our results for other countries. For this, we rely on data from the OECD about distributional consumption expenditures.[25] This data is

[24]More precisely, the correlation of food are 0.93, 0.94 and 0.91 for EPF V, EPF VI, and EPF VII, respectively. For services, those correlations are 0.83, 0.91 and 0.86.

[25]See the section of Annual National Accounts at https://stats.oecd.org.

FIGURE A.2: Consumption patterns across different waves of the Consumption Expenditure Survey

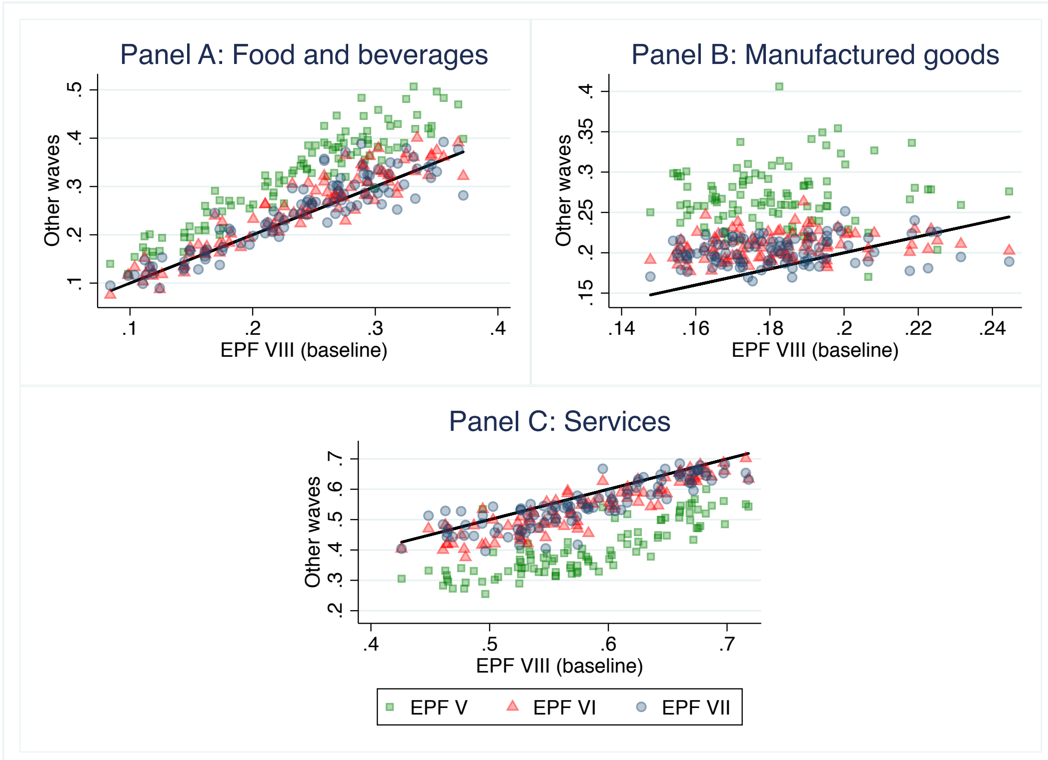

NOTES: This figure compares the share in consumption expenditures for households in each percentile of the income distribution between different waves of the Consumption Expenditure Survey (EPF). On every panel, each point corresponds to a percentile. Black solid line denotes 90 degree line.

an unbalanced panel for 13 countries in the period 1999-2019.[26] Different to the Chilean data, we only observe expenditures (in domestic currency) for the (equivalized) disposable income quintiles and not at the household level. To make progress, we define a dummy variable Q_k taking value equal to one for observations corresponding to quintile $k = 1, \ldots, 5$, and zero otherwise. As in the Chilean data, we observe expenditures for the 12 divisions in the Classification of Individual Consumption by Purpose (COICOP) (see Table A.1 for details) and we aggregate into our three-goods classification to have the expenditure share in food, manufactures and services. Those goods are indexed by j.

To study the expenditure patterns across the income distribution for different countries and periods, we run the following regression:

$$s_{jkit} = \alpha + \sum_{k=2}^{5} \beta_k Q_k + \delta X_{it} + \varepsilon_{jit},$$

for every $j = 1, \ldots, 3$. In the previous specification, the dependent variable s_{jkit} denotes the expenditure

[26]The list of countries is: Australia, Canada, Czech Republic, France, Ireland, Israel, Mexico, the Netherlands, New Zealand, Slovenia, Sweden, United Kingdom, and the United States.

share of quintile k group of country i-period t in good j, and X_{it} denote controls that vary across country and time. Because the previous specification excludes the effect of the first quintile, the coefficients of interest, β_k denote expenditure shares in goods j relative to the lowest share of the population.

Table A.2 presents the results, where we control for GDP per capita as well as country and year fixed effects. Column (1) presents the results for expenditures on food. Similar to the Chilean case, our results show a decreasing pattern across the income distribution. For example, households in the second quintile spend 1.3% less than the first quintile, while the highest quintile spends 7.9% less. We observe a similar pattern (and magnitudes) in column (2) for the case of manufactures. Finally, column (3) shows the results for services. As in the case of Chile, expenditures in services are increasing in income. The second quintile spends 2.8% more than the first quintile. For the highest quintile this figure is five times larger, with an expenditure share 13.4% higher.

Also note that the effect of aggregate income, measured by GDP per capita, is consistent with the disaggregate patterns: a higher level of income is associated with a lower expenditure share in food and a higher expenditure share in services.

TABLE A.2: Consumption expenditures in OECD countries

	Food	Manufactures	Services
	(1)	(2)	(3)
Quintile 2	-0.013***	-0.015**	0.028***
	(0.004)	(0.005)	(0.003)
Quintile 3	-0.028***	-0.031***	0.060***
	(0.007)	(0.009)	(0.006)
Quintile 4	-0.047***	-0.045***	0.092***
	(0.010)	(0.013)	(0.011)
Quintile 5	-0.079***	-0.055**	0.134***
	(0.017)	(0.019)	(0.019)
Log (GDP pc)	-0.108***	-0.075	0.183***
	(0.007)	(0.055)	(0.058)
Observations	370	370	370
Adj R2-within	0.649	0.471	0.824

NOTES: This table presents estimates of the relationship between consumption expenditures and quintiles of the income distribution. All regressions includede country and year fixed effects. Standard errors clustered at the country and year level. ***, ** and * denote statistical significance at the 1, 5 and 10% levels, respectively.

A.3 Data Validation

While the focus of this paper is in the response of households to income shocks and their aggregate consequences, Sections 2.3 and 2.4 rely on data at the municipality level. How well this data approximate household heterogeneity and inequality? At the same time, our microdata collects information about expenditures using transaction level data. How well this data approximates aggregate expenditures?

Figure A.3 shows that the income classification at the municipality level tracks quite well the mode, median and average income of individuals within a municipality. This is, the "representative" agent at the municipality level tracks the income strongly correlates with the level of income of the municipality itself. Therefore, studying the response of expenditure shares gives valuable information about the response of expenditure shares across the income distribution of individuals.

FIGURE A.3: Income Classification at the Municipality level vs Individual Level

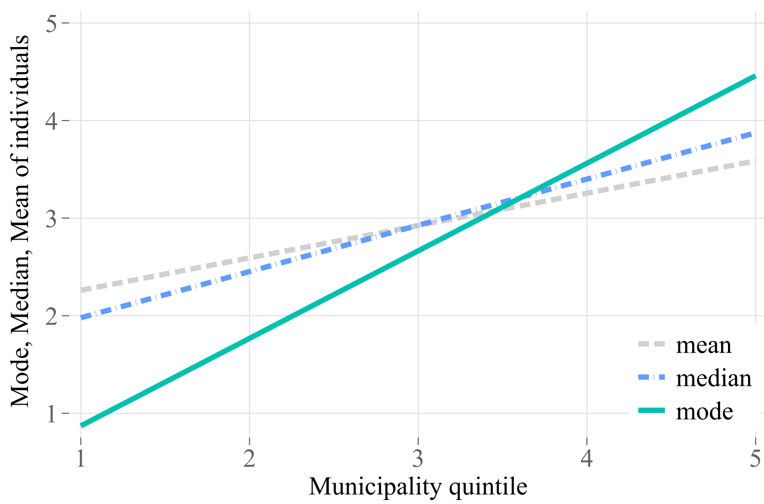

NOTES: This figure presents the relationship between income classification at the municipality level and the mode, median and mean of income across individuals within a municipality.

Regarding the validity of our data to capture the evolution of consumption, Figure A.4a compares aggregate consumption (taken from the National Accounts produced by the Central Bank of Chile) and aggregate expenditures coming from our Transbank microdata. As can be seen, our data tracks remarkably well the aggregate expenditure patterns of consumption. Moreover, this not only holds in the aggregate, but at the sectoral level, as shown in Figures A.4b and A.4c. The correlation of our transaction data are close to one even in the periods before COVID.

FIGURE A.4: Comparison of Transbank microdata with National Accounts

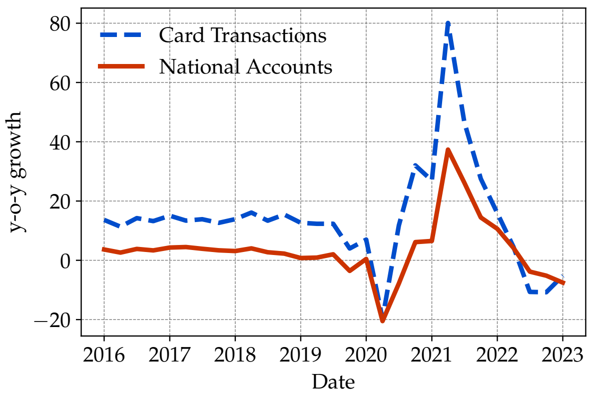

(a) Aggregate Expenditures

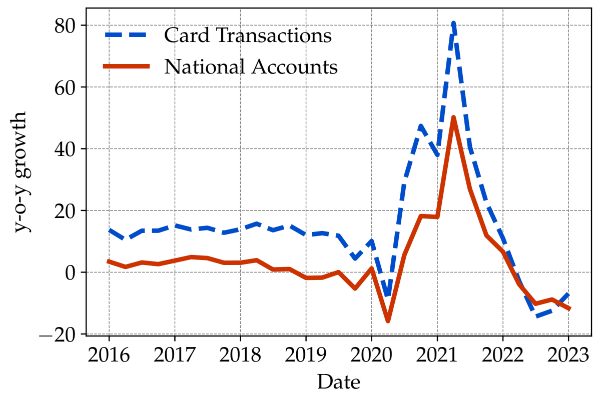

(b) Expenditures in Goods

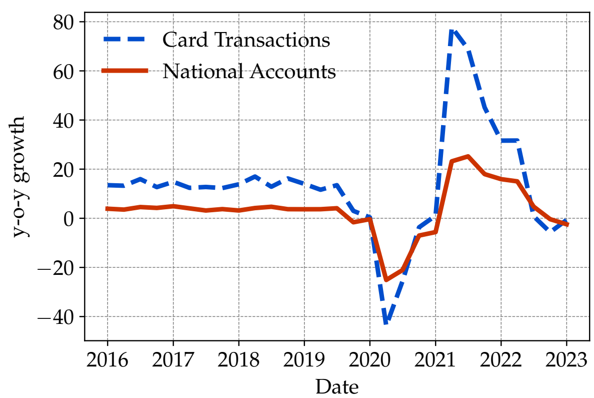

(c) Expenditures in Services

NOTES: This figure compares the evolution of sectoral consumption in Chile with sectoral expenditures coming from Transbank microdata. Panel a compares aggregate expenditures, while Panels b and c compares expenditures in goods and services respectively.

A.4 Additional Time-varying Evidence

This section provides additional empirical evidence regarding the time-varying behavior of expenditures presented in Sections 2.3 and 2.4.

Figure A.5 presents the unconditional correlation between expenditure shares and the business cycle for all the available data, including the COVID-19 period. As can be seen, the qualitative message presented in Section 2.3 remains.

Figure A.6 presents the response of expenditure shares in our local projection like exercise, controlling for the mobility index at the municipality level. As can be seen, the qualitative message presented in Section 2.4 remains.

A.5 Cyclicality of Aggregate Expenditure Shares in the Full Sample

FIGURE A.5: Ciclicality of Aggregate Expenditure Shares. Full Sample.

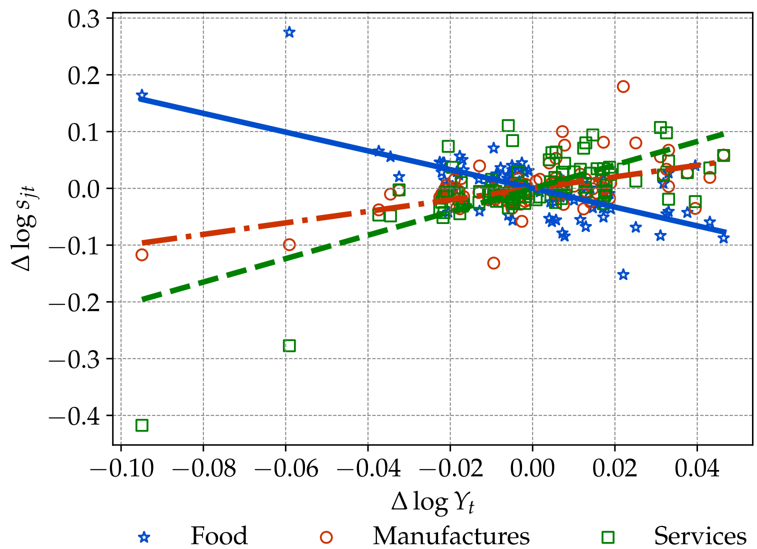

NOTES: This figure presents the unconditional correlation between the business cycle of sectoral expenditure shares and aggregate production index, including the COVID-19 period.

A.6 Robustness of Local Projections

FIGURE A.6: Response of Consumption Shares to Fiscal Transfers–Controlling for Mobility with 12 lags

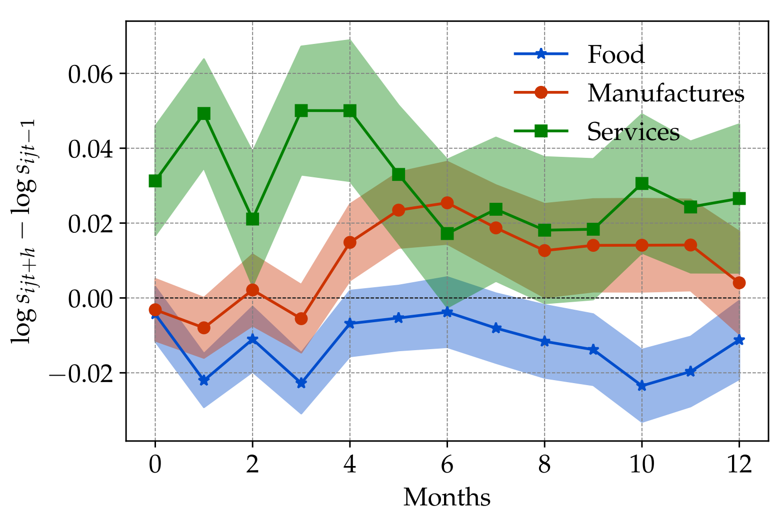

NOTES: This figure presents the dynamic response of expenditure shares to fiscal transfers (β_{jh} coefficients in. Eq. 3), controlling for the mobility index.

FIGURE A.7: Response of Consumption Shares to Fiscal Transfers Controlling for 8 Lags

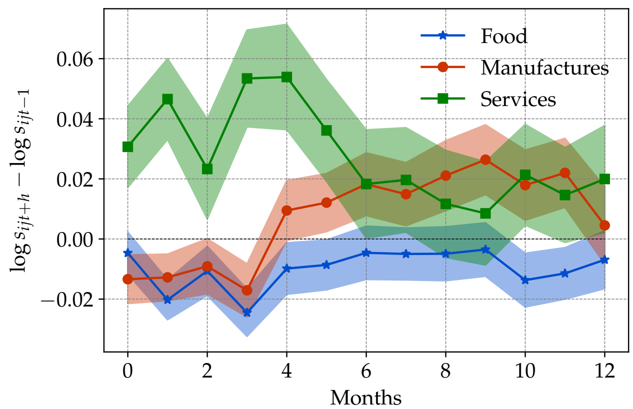

(a) Including Mobility Control

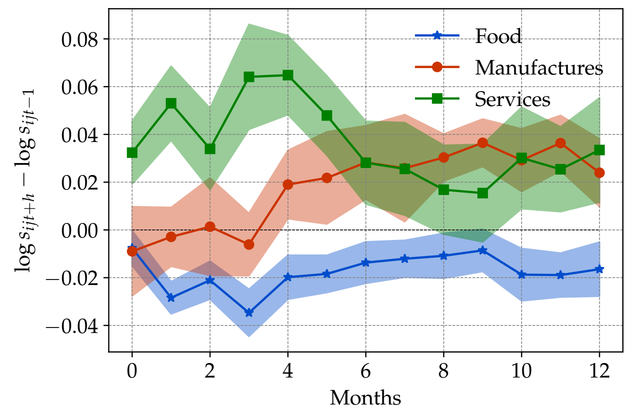

(b) Without Mobility Control

NOTES: This figure presents the dynamic response of expenditure shares to fiscal transfers (β_{jh} coefficients in. Eq. 3), controlling for the mobility index. Panel a controls for mobility while Panel b does not.

B Model Details

This section provides a self-contained description of the model presented in Section 3 that we implement in the quantitative section. In particular, we focus on the problem faced by households, labor unions and firms.

B.1 Households

There is a measure of one of households that differ in wealth b and productivity z. The b-dimension is the endogenous state, and the z-dimension is the exogenous time-varying idiosyncratic shock. We identify the type of households by i, which is associated to the tuple (b, z).

Households derive utility from consumption and disutility from labor. We assume there are a continuum of $g \in (0, 1)$ labor tasks that each household (b, z) can execute. Hence, household i maximizes its lifetime utility, time-discounted at a factor $0 < \beta < 1$, given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(c_{it}, \{n_{it}^g\}_{g=0}^1\right) \quad \text{with} \quad u(\{n_{it}^g\}_{j=0}^1) = \frac{(c_{it}-1)^{1-\gamma}}{1-\gamma} - \chi \frac{\int_0^1 (n_{it}^g)^{1+\varphi} dg}{1+\varphi},$$

where γ is the inverse of the intertemporal elasticity of substitution, χ is the parameter of the disutility of labor, and φ is the inverse of the Frisch elasticity of the labor supply, c_{it} is total consumption and n_{it}^g is hours supplied by workers from household i to the task g. Workers do not choose their labor supply directly due to labor market frictions. The labor supply is determined by a union that represents labor supplied by households to task g, in order to maximize the average welfare of all households providing those services. We first show the individual problem and in the next subsection, the wage-setting process faced by the union. The implication of that assumption is that $n^g_{it} = n_t, \ \forall \ i$. Consequently, nominal wages are common across households and equal to W_t.

Households maximize their lifetime utility subject to the following budget constraint

$$E_{it} + B_{it+1} = (1 + i_t)B_{it} + W_t n_t z_t + T_{it} - \tau_{it} + D_{it},$$

where $E_{it} \equiv P_{it}c_{it}$ denotes total expenditures that are also given by

$$E_{it} = \sum_{j=1}^{J} P_{jt}c_{ijt}.$$

In our setting, households consume a bundle of J goods facing a (common) price P_{jt} and consuming a quantity c_{ijt}. We denote by P_{it} the price index at the household level, while c_{it} is the total consumption of household i. As we explain below, our non-homothetic preferences generate household-level price indices, which imply that each household face a different costs of living, despite observing the same sectoral prices.

Households save in a liquid and risk-free asset B_{it} that returns a nominal interest rate i_t, receive income from labor $W_t n_t z_t$, where $W_t n_t$ is a "common" labor income and z_t is an individual idiosyncratic shock which generates a non-degenerate distribution of income. Households receive a fiscal transfer given by T_{it} and pay lump-sum taxes τ_{it}. Both fiscal instruments can have either a progressive or a flat profile. Finally, there are firms' dividends D_{it} that can also be distributed unevenly among the different households.

As described in the main text, the households' optimization problem can be split into an intratemporal problem, which describes how households distribute consumption among the different goods (which in our setup are nontrivial), and the intertemporal problem, which is also affected by the allocation of resources to the different goods.

Intratemporal Problem Households derive utility from the consumption of the J different goods in the economy (for our quantitative exercises we can think about the consumption basket as composed of food, manufactures, and services). The consumption aggregator of every household is denoted by c_{it} and takes the form of an implicitly additive non-homothetic CES function. It defines total consumption in period t by

$$1 = \sum_{j=1}^{J} \left(\omega_j \left(c_{it}\right)^{\epsilon_j}\right)^{\frac{1}{\sigma}} \left(c_{ijt}\right)^{\frac{\sigma-1}{\sigma}},$$

where c_{ijt} denotes the consumption of good j by household i in period t, ω_j is a taste parameter for good j (common across households), σ is the constant elasticity of substitution between sectoral goods, and ϵ_j is the constant elasticity of consumption of sectoral good j with respect to the consumption index c_{it} that allows preferences to be non-homothetic. This latter parameter is also understood as the income elasticity with respect to individual consumption goods.

Given a level of total expenditures $E_{it} = P_{it}c_{it}$, the intratemporal cost minimization optimization problem derives the following conditions

$$\begin{aligned} c_{ijt} &= \omega_j \left(\frac{P_{jt}}{P_{it}}\right)^{-\sigma} (c_{it})^{\epsilon_j+\sigma} \\ s_{ijt} \equiv \frac{P_{jt}c_{ijt}}{E_{it}} &= \omega_j \left(\frac{P_{jt}}{P_{it}}\right)^{1-\sigma} (c_{it})^{\epsilon_j-(1-\sigma)}, \end{aligned}$$

where s_{ijt} is the expenditure share of household i in good j. The household-specific CPI corresponds to the price index that equalizes $P_{it}c_{it} = \sum_{j=1}^{J} P_{jt}c_{ijt}$ and is defined by

$$P_{it} = \left[\sum_{j=1}^{J} (\omega_j P_{jt}^{1-\sigma})^{\vartheta_j} (s_{ijt}E_{it}^{1-\sigma})^{1-\vartheta_j}\right]^{\frac{1}{1-\sigma}},$$

with $\vartheta_j \equiv (1-\sigma)/\epsilon_j$. From these expressions is clear that both the expenditure share and the CPI of each household (characterized by different levels of income and access to financial markets) depend on the level of consumption/expenditures in every period. Note again that in the case of homothetic preferences ($\epsilon = 1-\sigma$), the expenditure shares do not depend on the level of consumption, and the CPI is common across households because it only depends on observed prices and not on the level of consumption itself (i.e., $\vartheta_j = 1$ holds). At the same time, with non-homothetic preferences, the demand for each good nonlinearly depends on total consumption through the good-specific income elasticity, ϵ_j. For future reference, denote $\bar{\epsilon}_{it} \equiv \sum_{j=1}^{J} s_{ijt}\epsilon_j$ as the average (expenditure-weighted) income elasticity. As we will see next, this object plays a crucial role in determining *intertemporal* consumption.

Intertemporal Problem. The intratemporal problem of the household derives the following Euler equation

$$u'(c(b,z)) = \beta\mathbb{E}_t\left[\frac{1+r}{1+\pi(b',z')}\frac{\bar{\epsilon}(b,z)}{\bar{\epsilon}(b',z')}u'(c(b',z'))\right] + \beta\mu(b',z')\frac{p(b,z)\bar{\epsilon}(b,z)}{1-\sigma}.$$

Households' Distribution. The above consumption-savings problem generates a distribution of households in the space $\mathcal{B}\times\mathcal{Z}$, where $\mathcal{B}$ is determined by the borrowing constraint and $\mathcal{Z}$ by the stochastic process governing z. We denote the distribution of households by $\Psi(b,z) = \Psi(i)$, which satisfies

$\int\int \Psi(i)di = 1$. Due to the recursive formulation of the problem, and given the policy functions of households, there is an operator F that maps $\Psi(i)$ onto $\Psi'(i)$

$$\Psi'(i) = F\left(\Psi(i)\right).$$

The mapping $F(\cdot)$ is a key object of our analysis because it keeps track of the distribution of households both in the steady state and along the transition path of the economy.

B.2 Workers' Union

We assume that for each task g, there is a union that decides wages w_t^g and the labor supplied n_t^g. In this setting, unions have market power as workers' tasks are in monopolistic competition. The union aggregates individual labor such that $n_t^g = \int\int n_t^g(b,z)dbdz$. Then, we assume there is a Dixit-Stiglitz aggregator that determines aggregate labor, given by

$$N_t = \left(\int_0^1 \left(n_t^g\right)^{\frac{\varepsilon-1}{\varepsilon}} dg\right)^{\frac{\varepsilon}{\varepsilon-1}},$$

where ε is the elasticity of the demand for labor tasks, which is also a measure of the market power of the union. The Dixit-Stiglitz aggregator gives rise to the following demand for each task g:

$$n_t^g = \left(\frac{w_t^g}{w_t}\right)^{-\varepsilon} N_t. \tag{B.1}$$

We assume nominal wages are sticky and their changes are subject to the following Rotemberg adjustment costs that are measured in utility units

$$\Gamma\left(\frac{W_t^g}{W_{t-1}^g} - 1\right) = \frac{\theta_w}{2}\left(\frac{W_t^g}{W_{t-1}^g} - 1\right)^2, \tag{B.2}$$

where θ_w is the nominal wage adjustment cost parameter. Then, the problem of the union is to choose the optimal labor, the nominal wage, and the wage inflation rate by solving:

$$\max_{n_{it}^g, W_{it}^g, \pi_{wt}^g} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int\int \left\{U\left(c_{it}\right) - v\left(n_{it}^g\right)\right\} \Psi_t(i)di - \Gamma^s\left(\frac{W_t^g}{W_{t-1}^g} - 1\right)\right], \tag{B.3}$$

subject to (B.1), and given that wage inflation is defined as $\pi_{wt}^g = \frac{W_t^g - W_{t-1}^g}{W_{st-1}^g}$. After imposing symmetry, this maximization problem leads to:

$$(\pi_{wt}+1)\pi_{wt} = \frac{\varepsilon_w}{\theta_w} n_t \int\int \left\{v'(n_t) - \frac{\varepsilon_w - 1}{\varepsilon_w} U'(c_{it}) \frac{W_t}{P_{it}} \frac{1-\sigma}{\bar{\epsilon}_{it}}\right\} \Psi_t(i)di + \beta\theta_w(\pi_{wt+1}+1)\pi_{wt+1},$$

which is the New Keynesian Wage Phillips Curve (NKWPC). The NKWPC relates the nominal wage inflation with hours worked and worker's preferences; it is a version of the wage Phillips curve described by Galí (2013) adapted to heterogeneity, Rotemberg adjustment costs, and non-homoteticities. Due to labor market frictions, all workers supply N_t hours at a real wage w_t. The NKWPC shows that non-homoteticities also enter the labor supply schedule through the average income elasticity of the different households $\overline{\epsilon}_{it}$. This implies that not only fluctuations in total consumption affect the labor supply but the composition of expenditures represented by this elasticity. To be more clear, let us write the labor supply implied by these expressions and our assumed preferences:

$$\psi N_t^{\varphi} = \frac{1}{\mathcal{M}_{wt}} W_t \int\int \left(\frac{(1-\sigma)}{\overline{\epsilon}_t(b,z)\ P_{it}(c_{it})^{\gamma}} \right) \Psi_t(i) di. \tag{B.4}$$

Equation (B.4) is the labor schedule in this economy. It has implicit the fact that unions have market power and wages are rigid through the dependence on $\mathcal{M}_{wt}$ which is the wage markup that workers get. We can rewrite this expression by multiplying by aggregate consumption (to the power of γ) to obtain:

$$\mathcal{M}_{wt}\psi N_t^{\varphi} C_t^{\gamma} = \frac{W_t}{P_t} \int\int \left(\frac{(1-\sigma)}{\overline{\epsilon}_{it}} \frac{P_t}{P_{it}} \left(\frac{C_t}{C_{it}} \right)^{\gamma} \right) \Psi_t(i) di. \tag{B.5}$$

B.3 Firms

We assume there are J sectors composed of a final good producer and intermediate producers. Final goods are a composite of a measure one of intermediates that operate in monopolistic competition. All sectors demand workers, but at a different levels. We describe the setup and optimality conditions for a given sector, denoted by j.

Final Goods Producers. In a given sector j, a competitive representative firm produces a final good by aggregating a continuum of measure one intermediate good with the following production function

$$Y_{jt} = \left(\int_0^1 {y_{jt}^m}^{\frac{\varepsilon-1}{\varepsilon}} dm \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

In this setting, the final firm decides how to allocate its demand among the different intermediate goods. After cost minimization, the demand for each intermediate good m, and the sector j price index writes:

$$y_{jt}^m = \left(\frac{p_{jt}^m}{p_{jt}} \right)^{-\epsilon} Y_{jt}, \quad \text{and} \quad p_{jt} = \left(\int_0^1 (p_{jt}^m)^{1-\varepsilon} dm \right)^{\frac{1}{1-\varepsilon}}. \tag{B.6}$$

Intermediate Goods Producers: Labor Demands. Each intermediate good m in sector j is produced by a monopolistically competitive producer using labor n_{jt}^m according to the production function

$$y_{jt}^m = A_{jt} n_{jt}^{1-\alpha}.$$

Each intermediate producer hires workers at a nominal wage W_t. Therefore, the demand satisfies

$$W_t = mc_{jt}(1-\alpha)N_{jt}^{-\alpha},$$

where N_{jt} is hours worked in sector j and mc_{jt} is the nominal marginal cost of sector j. Moreover, we can define the price markup of sector j as $\mathcal{M}_{jt}^p = \frac{p_{jt}}{mc_{jt}}$, which is different from one due to monopolistic competition and fluctuates due to price rigidities.

Intermediate Goods Producers: Price Setting. Each intermediate producer chooses its price to maximize profits subject to Rotemberg (1982) price adjustment costs. These adjustment costs are quadratic in the rate of price change $\frac{p_{jt}^m}{p_{jt-1}^m} - 1$ and are expressed as a fraction of output Y_t

$$\Theta_{jt}^m = \frac{\theta_j^p}{2}\left(\frac{p_{jt}^m}{p_{jt-1}^m} - 1\right)^2 Y_t \tag{B.7}$$

Therefore, each intermediate producer chooses $\{p_{ft}^m\}_{t\geq 0}$ to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty}\left(\frac{1}{1+r_t}\right)\left\{\Pi_t(p_{jt}^m) - \Theta_t\right\},$$

with

$$\Pi_{jt}^m(p_{jt}^m) = \left(\frac{p_{jt}^m}{P_{jt}} - mc_{jt}\right)\left(\frac{p_{jt}^m}{P_{jt}}\right)^{-\varepsilon} y_{jt},$$

where we assume that $\left(\frac{1}{1+r_t}\right)$ is the relevant discount factor of the firm, and mc_{jt} is the marginal cost. Given the assumptions above, the inflation rate (after the intermediate firms' optimization) is determined by the following New Keynesian Phillips curve for sector j

$$(\pi_{jt} - \overline{\pi}_j)\pi_{jt} = \frac{\varepsilon}{\theta_j^p}\left(\frac{mc_{jt}}{p_{jt}} - \frac{\varepsilon - 1}{\varepsilon}\right) + \mathbb{E}_t\left[\left(\frac{1}{1+r_t}\right)(\pi_{jt+1} - \overline{\pi}_j)\pi_{jt+1}\frac{p_{jt+1}y_{jt+1}}{p_{jt}y_{jt}}\right],$$

with $\pi_{jt} = \frac{p_{jt}}{p_{jt-1}}\pi_t$. Intermediate firms generate each period an aggregate amount of profits given by

$$D_{jt} = (p_{jt} - mc_{jt})\,Y_{jt} - \frac{\theta_j}{2}\pi_{jt}^2 Y_{jt} + mc_{jt}\alpha Y_{jt}$$

$$D_{jt} = p_{jt}Y_{jt} - w_t N_{jt}.$$